arXiv:1609.05934v1 [cs.DM] 19 Sep 2016

A simple algorithm for sampling colourings of $G(n, d/n)$ up to Gibbs Uniqueness Threshold *

Charilaos Efthymiou†

Goethe University
Frankfurt am Main, 60325, Germany
E-mail: efthymiou@gmail.com

September 21, 2016

Abstract

Approximate random k-colouring of a graph G is a well studied problem in computer science and statistical physics. It amounts to constructing a k-colouring of G which is distributed close to *Gibbs distribution* in polynomial time. Here, we deal with the problem when the underlying graph is an instance of Erdős-Rényi random graph $G(n, d/n)$, where d is a sufficiently large constant.

We propose a novel efficient algorithm for approximate random k-colouring $G(n, d/n)$ for any $k \geq (1 + \epsilon)d$. To be more specific, with probability at least $1 - n^{-\Omega(1)}$ over the input instances $G(n, d/n)$ and for $k \geq (1+\epsilon)d$, the algorithm returns a k-colouring which is distributed within total variation distance $n^{-\Omega(1)}$ from the Gibbs distribution of the input graph instance.

The algorithm we propose is neither a MCMC one nor inspired by the message passing algorithms proposed by statistical physicists. Roughly the idea is as follows: Initially we remove sufficiently many edges of the input graph. This results in a "simple graph" which can be k-coloured randomly efficiently. The algorithm colours randomly this simple graph. Then it puts back the removed edges one by one. Every time a new edge is put back the algorithm updates the colouring of the graph so that the colouring remains random.

The performance of the algorithm depends heavily on certain spatial correlation decay properties of the Gibbs distribution.

Key words: Random colouring, sparse random graph, efficient algorithm.

AMS subject classifications: Primary 68R99, 68W25,68W20 Secondary: 82B44

*Parts of this work appeared in SODA 2012 [8] and ESA 2014 [9].

†This work is supported by Deutsche Forschungsgemeinschaft (DFG) grant EF 103/11

1 Introduction

Let $\boldsymbol{G} = G(n, d/n)$ denote the random graph on the vertex set $V(\boldsymbol{G}) = \{1, \ldots, n\}$ where each edge appears independently with probability d/n, for a sufficiently large fixed number $d > 0$.

Approximate random k-colouring of a graph G is a well studied problem. It amounts to constructing a k-colouring of G which is distributed close to *Gibbs distribution*, i.e. the uniform distribution over all the k-colourings of G, in polynomial time. Here, we consider the problem when the underlying graph is an instance of Erdős-Rényi random graph $\boldsymbol{G} = G(n, d/n)$. This problem is a rather natural one and it has gathered focus in computer science but also in statistical physics.

From a technical perspective, the main challenge is to deal with the so called *effect of high degree* vertices. That is, there is a relative large fluctuation on the degrees in $\boldsymbol{G}$. E.g. it is elementary to verify that the typical instances of $\boldsymbol{G}$ have maximum degree $\Theta\left(\frac{\log n}{\log\log n}\right)$, while in these instances more than $1 - e^{-O(d)}$ fraction of the vertices have degree in the interval $(1 \pm \epsilon)d$. Usually the bounds for sampling k-colourings w.r.t. k are expressed it terms of the *maximum degree* e.g. [18, 6, 10, 11, 15]. However, for $\boldsymbol{G}$ it is natural to have bounds for k expressed in terms of the *expected degree* d, rather than the maximum degree.

The related work on this problem can be divided into two strands. The first one is based on *Markov Chain Monte Carlo* (MCMC) approach. There, the goal is to prove that some appropriately defined Markov Chain[1] over the k-colourings of the input graph is rapidly mixing. The MCMC approach to the problem is well studied [7, 5, 17]. The most recent of these works, i.e. [7], shows that the well known Markov chain *Glauber block dynamics* has polynomial mixing time for typical instances of $\boldsymbol{G}$ as long as the number of colours $k \geq \frac{11}{2}d$. This is the lowest bound for k as far as MCMC sampling is concerned.

The second strand has been based on message passing algorithms such as *Belief propagation* [4], which are closely related to the (non-rigorous) statistical mechanics techniques for the analysis of the random graph colouring problem. These message passing algorithms aim to approximate (conditional) *marginals* of the Gibbs distribution at each vertex . Given the marginals, a colouring can be sampled by choosing a vertex v, assigning it a random colour i according to the marginal distribution, and repeating the procedure with the colour of v fixed to i. Of course, the challenge is to prove that the algorithm does indeed yield sufficiently good estimates of the marginals. In a similar spirit, and subsequently to this work, the authors of [21] propose an approximate random colouring algorithm for $\boldsymbol{G}$ which uses the so-called Weitz's computational tree approach, from [20], to compute Gibbs marginals for colorings. This algorithms requires at least $3d$ many colours for the running time to be polynomial, i.e. $O(n^s)$ for some $s = s(d) > 0$.

In this work we obtain a considerable improvement over the best previous results by presenting a novel algorithm that only requires $k = (1 + \epsilon)d$ colours. The new algorithm does not fall into any of the categories discussed above. Instead, it rests on the following approach: Given the input graph, first remove sufficiently many vertices such that the resulting graph has a "very simple" structure and it can be randomly k-coloured *efficiently*. Once we have a random colouring of this, simple, graph we start adding one by one all the edges we have removed in the first place. Each time we put back in the graph an edge we *update* the colouring so that the new graph remains (asymptotically) randomly coloured. Once the algorithm has rebuilt the initial graph it returns its colouring.

Perhaps the most challenging part of the algorithm is to *update* the colouring once we have added an extra edge. The problem can be formulated as follows. Consider two fixed graphs G and G' such that $V(G) = V(G')$ and $E(G') = E(G) \cup \{v, u\}$ for some $v, u \in V(G)$. Given X, a random k-colouring of G, we want to create *efficiently* a random k-colouring of the slightly more complex graph G'. It is easy to show that if the vertices v, u have different colour assignments under X, then X is a random k-

[1] e.g. Glauber dynamics

colouring of G'. The interesting case is when $X(v) = X(u)$. Then the algorithm should alter the colour assignment of at least one of the two vertices such that the resulting colouring is random conditional that the assignments of v and u are different. Here, we use an operation which we call "switching" so as to alter the colouring of only one of the two vertices. Roughly speaking, the switching chooses an appropriately large part of G, which contains only v. Then, it repermutes appropriately the colour classes in this part of G so as to get the updated colouring.

For presenting our results we use the notion of *total variation distance*, which is a measure of distance between distributions.

Definition 1 *For the distributions ν_a, ν_b on $[k]^V$, let $||\nu_a - \nu_b||$ denote their* total variation distance*, i.e.*

$$||\nu_a - \nu_b|| = \max_{\Omega' \subseteq [k]^V} |\nu_a(\Omega') - \nu_b(\Omega')|.$$

For $\Lambda \subseteq V$ let $||\nu_a - \nu_b||_\Lambda$ be the total variation distance between the projections of ν_a and ν_b on $[k]^\Lambda$.

Theorem 1 *Let $\epsilon > 0$ be a fixed number, let d be sufficiently large number and fixed $k \geq (1+\epsilon)d$. Consider $\boldsymbol{G} = G(n, d/n)$ and let μ the uniform distribution over the k-colouring of $\boldsymbol{G}$. Let $\hat{\mu}$ be the distribution of the colouring that is returned by our algorithm on input $\boldsymbol{G}$.*

Let $c = \frac{\epsilon}{80(1+\epsilon/4)\log d}$, with probability at least $1 - n^{-c}$ over the input instances $\boldsymbol{G}$ it holds that

$$||\mu - \hat{\mu}|| = O\left(n^{-c}\right). \tag{1}$$

The proof of Theorem 1 appears in Section 6.

The following theorem is for the time complexity of the algorithm, its proof appears in Section 6.

Theorem 2 *With probability at least $1 - 2n^{-2/3}$ over the input instances $\boldsymbol{G}$, the time complexity of the random colouring algorithm is $O(n^2)$.*

Whether the running time of the algorithm is polynomial or not, depends on certain structural properties of the input graph $\boldsymbol{G}$. Mainly, these properties require that the "short cycles" of $\boldsymbol{G}$ are disjoint. It will be trivial to distinguish the instances that can be coloured randomly efficiently by our algorithm from those that cannot, see in Section 6 for further details.

Remark 1 *The region of k for which our algorithm operates, coincides with what is conjectured to be the so-called "*Uniqueness phase*" of the k-colourings of $\boldsymbol{G}$, e.g. see [22].*

Remarks on the accuracy Typically, the approximation guarantees we get from algorithms as those in [7, 21] express the running time of the algorithm as a polynomial of the error in the output. The running time and the error of the algorithm here are independent, in the sense that the approximation guarantees do not improve by allowing the algorithm run more steps.

Notation Given some graph G, we let $V(G)$ and $E(G)$ denote the vertex sets and the edge set, respectively. Also, we let $\Omega_{G,k}$ be the set of proper k-colourings of G. We denote with small letters of the greek alphabet the colourings in $\Omega_{G,k}$, e.g. σ, η, τ. We use capital letters for the random variables which take values over the colourings e.g. X, Y, Z. We denote with $\sigma_v, X(v)$ the colour assignment of the vertex v under the colouring σ and X, respectively. Given some $\sigma \in \Omega_{G,k}$, for every $i \in [k]$ we let $\sigma^{-1}(i) \subseteq V(G)$ be the colour class of colour i under the colouring σ. Finally, for some integer $h > 0$, we let $[h] = \{1, \ldots, h\}$.

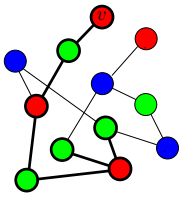

Figure 1: "Disagreement graph".

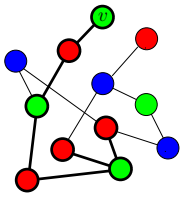

Figure 2: "switching".

2 Basic Description

So as to give a basic description of our algorithm, we need to introduce few notions. Consider a fixed graph G and let v be a vertex in $V(G)$. Let $c, q \in [k]$ be different with each other and let σ be a k-colouring of G such that $\sigma(v) = c$. We call *disagreement graph* $\boldsymbol{Q} = \boldsymbol{Q}(G, v, \sigma, q)$, the maximal, connected, induced subgraph of G such that $v \in V(\boldsymbol{Q})$, while $V(\boldsymbol{Q}) \subseteq \sigma^{-1}(c) \cup \sigma^{-1}(q)$.

Remark 2 *The concept of disagreement graph, in the graph theory literature is also known as Kempe Chain.*

In Figure 1, the disagreement graph $\boldsymbol{Q}(G, v, \sigma, \text{“}\texttt{green}\text{”})$ is the one with the fat lines. Note that σ specifies a two colouring for the vertices of $\boldsymbol{Q}(G, v, \sigma, \text{“}\texttt{green}\text{”})$.

Definition 2 *Consider G, v, σ and q as specified above, as well as the disagreement graph $\boldsymbol{Q} = \boldsymbol{Q}(G, v, \sigma, q)$. The "$q$-switching of σ" corresponds to the colouring of G which is derived by exchanging the assignments in the two colour classes in $\boldsymbol{Q}$.*

Figure 2 illustrates a switching of the colouring in Figure 1. That is, the colouring in Figure 2 differs from the one in Figure 1 in that we have exchanged the two colour classes of the subgraph with the fat lines. The q-switching of any proper colouring of G is always a proper colouring, too.

We proceed with a high level description of the algorithm. The input is $\boldsymbol{G} = G(n, d/n)$ and some integer $k \geq (1+\epsilon)d$. The algorithm is as follows:

Set up: We construct a sequence of graphs $G_0, \ldots, G_r$ such that G_r is identical to $\boldsymbol{G}$ and G_i is a subgraph of G_{i+1}. Each G_i is derived by deleting from G_{i+1} the edge $\{v_i, u_i\}$. This edge is chosen at random among those which do not belong to a *short cycle* of G_{i+1}. We call short, any cycle of length less than $(\log_d n)/9$. G_0 is the graph we get when there are no other edges to delete.

With probability $1 - n^{-\Omega(1)}$, over the instances of $\boldsymbol{G}$, the above process generates G_0 which is simple[2] enough that can be k-coloured randomly in polynomial time. If G_0 is not simple, the algorithm cannot proceed and abandons. Assuming that G_0 is simple, the algorithm proceeds as follows:

Update: Take a random colouring of G_0. Let Y_0 be that colouring. We get $Y_1, Y_2, \ldots, Y_r$, the colourings of $G_1, G_2, \ldots, G_r$, respectively, according to the following inductive rule: Given that G_i is coloured Y_i, so as to get Y_{i+1} we distinguish two cases

Case (a): Y_i (the colouring of G_i) assigns v_i and u_i different colours, i.e. $Y_i(v_i) \neq Y_i(u_i)$

Case (b): Y_i assigns v_i and u_i the same colour, i.e. $Y_i(v_i) = Y_i(u_i)$.

In the first case, we set $Y_{i+1} = Y_i$, i.e. G_{i+1} gets the same colouring as G_i. In the second case, we choose q uniformly at random from $[k] \backslash \{Y_i(v_i)\}$, i.e. among all the colours but $Y_i(v_i)$. Then, we set Y_{i+1} equal to the q-switching of Y_i. The q-switching is w.r.t. the graph G_i, the vertex v_i and the colouring Y_i. The algorithm repeats these steps for $i = 0, \ldots, r-1$. Then it outputs Y_r.

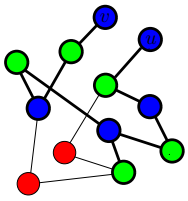

G

Figure 3: "Disagreement graph".

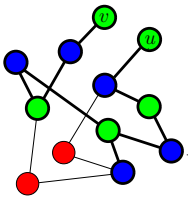

G

Figure 4: "switching".

One could remark that the switching does not necessarily provide a k-colouring where the assignments of v_i and u_i are different. That is, it may be that both vertices v_i, u_i belong to the disagreement graph in Y_i, e.g. Figure 3. Then, after the q-switching the colour assignments of v_i and u_i remain the same, e.g. Figure 4. It turns out that this situation is rare as long as $k = (1+\epsilon)d$. More specifically, with probability $1 - o(n^{-1})$, the q-switching of Y_i specifies different colour assignments for v_i, u_i.

The approximate nature of the algorithm amounts exactly to the fact that on some, rare, occasions the switching somehow fails. The error at the output of the algorithm (see Theorem 1) is closely related to the probability of the event that our algorithm encounters such failure when the input is a typical instance of $\boldsymbol{G}$.

Remark 3 *The lower bound we have for k depends exactly how well we can control these failures of switching. That is, for $k \leq d$ our analysis cannot guarantee that the switching fails only on rare occasions.*

3 The setting for the analysis of the algorithm.

Consider a fixed graph G and let v, u be two distinguished, non-adjacent, vertices.

Definition 3 (Good & Bad colourings) *Let σ be a proper k-colouring of G, for some $k > 0$. We call σ* bad colouring *w.r.t. the vertices v, u of G, if $\sigma_v = \sigma_u$. Otherwise, we call σ* good.

The idea that underlies the sampling algorithm, reduces the sampling problem to dealing with the following one.

Problem 1 *Given a* bad *random colouring of G, w.r.t. $\{v, u\}$, turn it to a* good *random colouring, in polynomial time.*

Consider two different $c, q \in [k]$ and let $\Omega_{c,c}$ and $\Omega_{q,c}$ be the set of colourings of G which assign the pair of vertices (v, u) colours (c, c) and (q, c), respectively. Our approach to Problem 1 relies on getting a mapping $H_{c,q} : \Omega_{c,c} \to \Omega_{q,c}$ such that the following holds:

A. If Z is uniformly random in $\Omega_{c,c}$, then $H_{c,q}(Z)$ is uniformly random in $\Omega_{q,c}$

B. The computation of $H_{c,q}(Z)$ can be accomplished in polynomial time.

It is straightforward that having such a mapping for every two $c, q \in [k]$, it is sufficient to solve Problem 1. In the following discussion our focus is on (the more challenging) $\mathbf{A}$. rather than $\mathbf{B}$.

An ideal (and to a great extent untrue) situation would have been if $\Omega_{c,c}$ and $\Omega_{q,c}$ admitted a bijection. Then for $\mathbf{A}$. it would suffice to use for $H_{c,q}$ a bijection between the two sets. Since this is not expected to hold in general, our approach is based on introducing an *approximate bijection* between the sets $\Omega_{c,c}$

[2] In our case, G_0 is considered simple if it is component structure is as follows: Each component is either an isolated vertex, or a simple isolated cycle. In Section 6 we describe how someone can get efficiently a random colouring of such a graph.

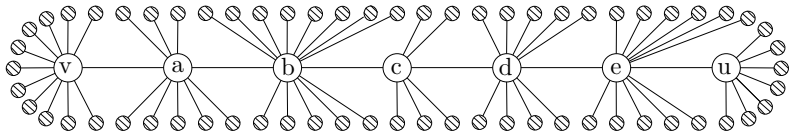

Figure 5: Boundary at distance 1 from the path.

and $\Omega_{q,c}$. That is, we consider a mapping which is a bijection between two sufficiently large subsets of $\Omega_{c,c}$ and $\Omega_{q,c}$, respectively. This would mean that if Z is uniformly random in $\Omega_{c,c}$ and $H_{c,q}(\cdot)$ an approximate bijection between $\Omega_{c,c}$ and $\Omega_{q,c}$, then $H_{c,q}(Z)$ is *approximately* uniformly random in $\Omega_{q,c}$.

To be more specific, we let $H_{c,q}$ represent the operation of q-switching over the colourings in $\Omega_{c,c}$, as we describe in Section 2. For such mapping, we can find appropriate $\Omega'_{c,c} \subseteq \Omega_{c,c}$ and $\Omega'_{q,c} \subseteq \Omega'_{q,c}$ such that $H_{c,q}$ is a bijection between the sets $\Omega_{c,c} \backslash \Omega'_{c,c}$ and $\Omega_{q,c} \backslash \Omega'_{q,c}$. We call *pathological* each colouring $\sigma \in \Omega'_{c,c} \cup \Omega'_{q,c}$. For the pathological colouring $\sigma \in \Omega'_{c,c}$ it holds that $H_{c,q}(\sigma) \notin \Omega_{q,c}$, while for $\sigma \in \Omega'_{q,c}$ it holds that $H^{-1}_{c,q}(\sigma) \notin \Omega_{c,c}$.

Remark 4 *There is a natural characterization for the pathological colourings $\sigma \in \Omega_{c,c}$. That is, σ is pathological if the disagreement graph $\boldsymbol{Q} = \boldsymbol{Q}(G, v, \sigma, q)$ contains both v, u.*

It turns out that, for Z being uniformly random in $\Omega_{c,c}$, $H_{c,q}(Z)$ is distributed within total variation distance $\max\left\{\frac{\Omega'_{c,c}}{\Omega_{c,c}}, \frac{\Omega'_{q,c}}{\Omega_{q,c}}\right\}$ from the uniform distribution over $\Omega_{q,c}$. That is, the error we introduce with the approximate bijection $H_{c,q}$ depends on the *relative number* of the pathological colorings in $\Omega_{c,c}$ and $\Omega_{q,c}$, respectively. A key ingredient of our analysis is to provide appropriate upper bounds for the two ratios $\Omega'_{c,c}/\Omega_{c,c}, \Omega'_{q,c}/\Omega_{q,c}$.

3.1 Bounding the Error - Spatial Mixing

As in the previous section, let G be fixed. For bounding the ratios $\Omega'_{c,c}/\Omega_{c,c}$ and $\Omega'_{q,c}/\Omega_{q,c}$, we treat both cases in the same way, so let us focus on bounding $\Omega'_{c,c}/\Omega_{c,c}$.

It is direct that $\Omega'_{c,c}/\Omega_{c,c}$ expresses the probability of getting a pathological colouring if we choose uniformly at random from $\Omega_{c,c}$. For this, consider the situation where we choose u.a.r. from $\Omega_{c,c}$. For every path P that connects v, u in the graph G, we let $\mathbf{I}_{\{P\}}$ be an indicator variable which is one if the vertices in the path P are coloured only with colours c, q in the random colouring and zero otherwise. Equivalently, $\mathbf{I}_{\{P\}} = 1$ if and only if P belongs to the graph of disagreement that is induced by the random colouring and the colour q. It holds that

$$\frac{\Omega'_{c,c}}{\Omega_{c,c}} = \Pr\left[\sum_P \mathbf{I}_{\{P\}} \geq 1\right] \leq \sum_P \Pr\left[\mathbf{I}_{\{P\}} = 1\right]. \tag{2}$$

The first equality follows from the fact that if both v, u belong to the disagreement graph, then there should be at least one path P such that $\mathbf{I}_{\{P\}} = 1$. The last inequality follows from the union bound.

Remark 5 *The above inequality bounds the relative number of pathological colourings in $\Omega_{c,c}$ (resp. in $\Omega_{q,c}$) with the expected number of paths from v to u which are coloured with c, q under a colouring which is chosen at random from $\Omega_{c,c}$ (resp. $\Omega_{q,c}$).*

In general, computing $\Pr[\mathbf{I}_{\{P\}} = 1]$ exactly is a formidable task to accomplish due to the complex structure we typically have in the underlying graph. For this reason we reside on computing upper bounds of this probability term.

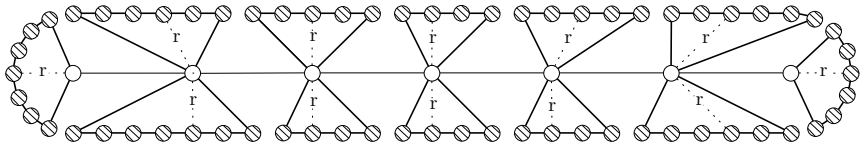

Figure 6: Boundary at distance r from the path

In [8] we used the idea of the so-called "Disagreement percolation" from [3]. The setting of this approach is illustrated in Figure 5, for the path $P = (v, a, b, c, d, e, u)$. The lined vertices are exactly these which are adjacent to the path. So as to bound the probability that the path P is coloured with c, q, we assume a worst case boundary colouring for the lined vertices. Given the fixed colourings at the boundary, we take a random colouring of the uncoloured vertices in P, conditional v, u are assigned c, and estimate the probability that P is coloured exclusively with c, q.

Remark 6 *The choice of the boundary above, is worst case in the sense that it maximizes the probability that* $\mathbf{I}_{\{P\}} = 1$.

It turns out that considering the worst case boundary condition next to the path P is a too pessimistic assumption. There is an improvement once we adopt a less restrictive approach. The new approach is illustrated in Figure 6. Roughly speaking, we consider a worst case boundary condition at the vertices around P which are at graph distance r, for $r \gg 1$. The boundary condition gives rise to Gibbs distribution over the k-colourings of the subgraph confined by the boundary vertices. In particular, we argue about the *spatial mixing* properties of the Gibbs distributions in the confined graph. We show that the colouring[3] of the distant vertices does not bias the distribution of the colour assignment of the vertices in P by too much.

The above approach is well motivated when we consider $G(n, d/n)$. For such graph, typically, around most of the vertices in P we have a tree-like neighbourhood of maximum degree very close to the expected degree d. This gives rise to study correlation decay for random colourings of a tree with maximum degree Δ, for $\Delta \approx d$. Our spatial mixing results build on the work of Jonasson [12].

From fixed graph to random graph. When the underlying graph G is fixed, we bound $\Omega'_{c,c}/\Omega_{c,c}$ (resp. $\Omega'_{q,c}/\Omega_{q,c}$) by using the expected number of paths between v and u that are coloured c, q in a colouring chosen uniformly at random from $\Omega_{c,c}$ (resp. $\Omega_{q,c}$). That is, we need to argue on the randomness of the k-colourings of G.

In our analysis, we deal with cases where the underlying graph is random. Then, we have an extra level of randomness to deal with, that of the graph instance. That is, we take an instance of the graph and then, given the graph, we consider a random colouring of this graph instance. Even in this setting, we compute the expected number number of paths between v and u that are coloured c, q, however, the expectation is w.r.t. to the randomness of both the graph and its colouring. A result which is central in our analysis is the following one.

Theorem 3 *Let $\epsilon > 0$, let $d > 0$ be sufficiently large and let fixed $k \geq (1 + \epsilon)d$. Consider $\boldsymbol{G} = G(n, d/n)$. Let the graph $\boldsymbol{H}$ be such that $V(\boldsymbol{H}) = V(\boldsymbol{G})$ and $E(\boldsymbol{H}) \subseteq E(\boldsymbol{G})$. For any two $c, q \in [k]$, different with each other, any non-negative integer $\ell \leq \log^2 n$ and a permutation $P = (w_0, \ldots, w_\ell)$ of vertices in $V(\boldsymbol{H})$ the following is true:*

[3] any colouring

Let X be a random k-colouring of $\boldsymbol{H}$ conditional than $X(w_0) = c$. Let $\mathbf{I}_{\{P\}} = 1$, if P is a path in $\boldsymbol{H}$ and $X(w_i) \in \{c, q\}$, for every $j = 1, \ldots, \ell$. Otherwise $\mathbf{I}_{\{P\}} = 0$. It holds that

$$\Pr[\mathbf{I}_{\{P\}} = 1] \leq 2[(1 + \epsilon/4)n]^{-\ell}. \tag{3}$$

The proof of the theorem appears in Section 9.

Remark 7 *In* (3) *the probability term is w.r.t. both the randomness of $\boldsymbol{H}$ and the colouring X.*

The above theorem implies that for $k \geq (1 + \epsilon)d$, in a random k-colouring of $\boldsymbol{G}$, typically, there are not long paths coloured with only two colours. Furthermore, this property is monotone in the graph structure. That is, it holds even though if we remove an arbitrary number of edges from $\boldsymbol{G}$ (and get $\boldsymbol{H}$). The monotonicity property follows from the fact that we can extend in a natural way the Gibbs uniqueness condition in [12] from Δ regular trees to trees of maximum degree Δ.

4 Updating Colourings

In this section, we describe the process that the random colouring algorithm uses to update the colourings, we call it `Update`. For the sake of clarity in this section we assume a fixed graph G and we distinguish two vertices $v, u \in V(G)$. We take k sufficiently large so that G is k-colourable.

Definition 4 (Disagreement graph) *For any $\sigma \in \Omega_{G,k}$ and $q \in [k]\backslash\{\sigma_v\}$ we let the disagreement graph $\boldsymbol{Q} = \boldsymbol{Q}(G, v, \sigma, q)$ be the maximal induced subgraph of G such that*

$$V(\boldsymbol{Q}) = \left\{ x \in V(G) \;\middle|\; \begin{array}{l} \exists \text{ path } w_1, \ldots, w_\ell, \text{ in } G \text{ such that:} \\ w_1 = v, w_\ell = x, \sigma(w_j) \in \{\sigma_v, q\}, \forall j \in [\ell] \end{array} \right\}.$$

Next, we provide the pseudo-code of the operation `Switching`, presented in Section 2.

Switching

Input: G, v, σ and $q \in [k]\backslash\{\sigma_v\}$
  `set` $c = \sigma_v$
  `set` $\boldsymbol{Q} = \boldsymbol{Q}(G, v, \sigma, q)$
  `set` $\tau(V(G)\backslash V(\boldsymbol{Q})) = \sigma(V(G)\backslash V(\boldsymbol{Q}))$ /* Everything outside $\boldsymbol{Q}$ keeps its initial colouring*/
  `for` $w \in V(\boldsymbol{Q}) \cap \sigma^{-1}(c)$ `do`
    `set` $\tau(w) = q$
  `for` $w \in V(\boldsymbol{Q}) \cap \sigma^{-1}(q)$ `do`
    `set` $\tau(w) = c$
Output: τ

`Switching` has the following property, whose proof is easy to derive.

Lemma 1 *If $\tau =$ `Switching`(G, v, σ, q), where $\sigma \in \Omega_{G,k}$ and $q \neq \sigma(v)$, then $\tau \in \Omega_{G,k}$.*

The proof of Lemma 1, is quite straightforward and appears in Section 13.1.

As far the time complexity of `Switching` is regarded we have the following lemma, whose proof appears in Section 13.2.

Lemma 2 *For every $v \in V(G)$, any $\sigma \in \Omega_{G,k}$, $q \in [k]\backslash\{\sigma_v\}$ the time complexity of computing* `Switching`(G, v, σ, q) *is $O(|E(G)|)$.*

In what follows, we have the pseudo-code for `Update`.

Update

Input: $G, v, u, \sigma \in \Omega_{G,k}$
`if` σ is a *good* colouring w.r.t. v, u, `then`
`set` $\tau = \sigma$
`else do`
`choose` q u.a.r. from $[k]\backslash\{\sigma_v\}$
`set` $\tau =$ `Switching`(G, v, σ, q)
Output: τ

To this end, we need argue about the time complexity and the accuracy of `Update`. As far as the time complexity is regarded we have the following theorem.

Theorem 4 *For any $v, u \in V$, $\sigma \in \Omega_{G,k}$ and $q \in [k]\backslash\{\sigma_v\}$, the time complexity of* `Update`(G, v, u, σ, k) *is $O(|E(G)|)$.*

Theorem 4 follows as a corollary of Lemma 2, once we note that the execution time of `Update` is dominated by the calls of `Switching`.

So as to study the accuracy of `Update` we introduce the following concepts. For any two different colours c, q we let $S_q(c, c) \subseteq \Omega(c, c)$ and $S_c(q, c) \subseteq \Omega(q, c)$ be defined as follows: The set $S_q(c, c)$ (resp. $S_c(q, c)$) contains every $\sigma \in \Omega(c, c)$ (resp. $\sigma \in \Omega(q, c)$) such that there is no path between v and u which is coloured only with the colours c, q, by σ.

Definition 5 *Let $\alpha = \alpha_{G,k} \in [0, 1]$ be the minimum number such that the following holds: For every pair of different colours $c, q \in [k]$ the sets $S_q(c, c)$ and $S_c(q, c)$ contain all but an α-fraction of colourings of $\Omega(c, c)$ and $\Omega(q, c)$, respectively.*

In general the value of α depends on the underlying graph G and k. The quantity α is an upper bound on the relative size of pathological colourings in each set $\Omega(c, c')$.

Theorem 5 *Let ν be the uniform distribution over the k-colourings of G which are* good*, w.r.t. v, u. Let, also, ν' be the distribution of the output of* `Update` *when the input colouring is distributed uniformly at random over the k-colourings of G. Letting α be as in Definition 5, it holds that*

$$||\nu - \nu'|| \leq \alpha.$$

The proof of Theorem 5 appears in Section 12.

5 Random Colouring Algorithm

In this section, we study the time complexity and the accuracy of the random colouring algorithm. For the sake of definitiveness we assume the input graph G to be fixed and is such that G is k-colourable. Given the input graph G, the algorithm creates the sequence of subgraphs $G_0, \ldots, G_r$. The variable Y_i denotes the k-colouring that the algorithm assigns to the graph G_i. G_i is derived by deleting from G_{i+1} an edge which we call $\{v_i, u_i\}$.

As we consider a general graph G, in the pseudo-code that follows, we do not specify exactly how do we get G_i from G_{i+1}, i.e. what is $\{v_i, u_i\}$. Also, we do not specify how do we get Y_0, the random colouring of G_0. We get specific on these two matters only when we consider $G(n, d/n)$ at the input, see Section 6.

The pseudo-code for the algorithm is as follows:

Random Colouring Algorithm

Input: G, k

`compute` $G_0, G_1 \ldots, G_r$

`compute` Y_0 /* *Get a random k-colouring of* G_0*/*

`for` $0 \leq i \leq r-1$ `do`

`set` Y_{i+1} the output of `Update`(G_i, v_i, u_i, Y_i, k)

Output: Y_r

Using Theorem 4 and noting that $r \leq |E(G)|$, we get the following result.

Theorem 6 *Let $T(n)$ be the time complexity for k-colouring randomly G_0. Then, the random colouring algorithm has time complexity $O(|E(G)|^2 + T(n))$.*

Next, we investigate the accuracy of the algorithm. For any $c, q \in [k]$ we let $\Omega_i(c, q)$ be the set of colourings of G_i which assign the colours c and q to the vertices v_i and u_i, respectively. Furthermore, for two different colours $c, q \in [k]$, let $S^i_q(c, c) \subseteq \Omega_i(c, c)$ and $S^i_c(q, c) \subseteq \Omega_i(q, c)$ be defined as follows: The set $S^i_q(c, c)$ (resp. $S^i_c(q, c)$) contains every $\sigma \in \Omega_i(c, c)$ (resp. $\sigma \in \Omega_i(q, c)$) such that there is no path between v_i and u_i (in G_i) which is coloured by σ using the colours c, q, only.

Definition 6 *For every $i = 0, \ldots, r-1$, let $\alpha_i \in [0, 1]$ be the minimum number such that the following holds: For any pair of different colours c, q the sets $S^i_q(c, c)$ and $S^i_c(q, c)$ contain all but an α_i-fraction of the colourings in $\Omega_i(c, c)$ and $\Omega_i(q, c)$, respectively.*

Clearly the quantities α_i depend on G_i and k.

Theorem 7 *Let μ be the uniform distribution over the k-colourings of the input graph G. Let $\hat{\mu}$ be the distribution of the colourings at the output of the algorithm. It holds that*

$$||\mu - \hat{\mu}|| \leq \sum_{i=0}^{r-1} \alpha_i,$$

where α_i is from Definition 6 and r is the number of terms of the sequence $G_0, G_1, \ldots, G_r$.

The proof of Theorem 7 appears in Section 13.3.

6 Random Colouring $G(n, d/n)$

In this section, we focus on the case where the input of `Random Colouring Algorithm` is $\boldsymbol{G} = G(n, d/n)$. This study leads to the proof of Theorems 1 and 2.

We start by describing how do we get $G_0, \ldots, G_r$ from $\boldsymbol{G}$. Let $\mathcal{E}(\boldsymbol{G}) \subseteq E(\boldsymbol{G})$ contain exactly every edge $e \in E(\boldsymbol{G})$ such that the shortest simple cycle that contains e is of length greater than $(\log_d n)/9$.

Computing $G_0, \ldots, G_r$: The sequence $G_0, \ldots, G_r$ is constructed as follows: Set $r = |\mathcal{E}| + 1$. We set $G_r = \boldsymbol{G}$. Given G_i we get G_{i-1} by removing a randomly chosen edge of G_i which also belongs to $\mathcal{E}(\boldsymbol{G})$, for $i = 1, \ldots, r$. G_0 contains only the edges of the initial graph which do not belong to $\mathcal{E}(\boldsymbol{G})$.

Perhaps it is interesting to describe what motivates the above construction of the sequence $G_0, \ldots, G_r$. Since each α_i depends on G_i, we construct the sequence so as to have $\sum_i \alpha_i$, as small as possible. The smaller the probability the algorithm encounters a disagreement graph which includes both v_i, u_i the smaller α_is get. Choosing v_i and u_i to be at large distance reduces the probability that the disagreement graph includes both of them, consequently, α_i gets smaller. Our choice of sequence forces v_i and u_i to

be at distance greater than $(\log_d n)/9$ with each other. To a certain extent, this allows to control the error of the algorithm, i.e. $\sum_i \alpha_i$.

Given the sequence $G_0, \ldots, G_r$, the next step is to argue on how can we get a random k-colouring of G_0, efficiently. Our arguments rely on the fact that typically G_0 has a very simple structure, i.e. we use the following result.

Lemma 3 *For $d > 0$, let $\mathfrak{S}_{n,d}$ be the set of all graph on n vertices such that their component structure is as follows: Each component is either the trivial[4], or it is a simple isolated cycle [5] of maximum length $(\log_d n)/9$. Consider $\boldsymbol{G}$ and the sequence $G_0, \ldots G_r$ created as we described above. It holds that*

$$\Pr[G_0 \in \mathfrak{S}_{n,d}] \geq 1 - n^{-2/3}.$$

The proof of Lemma 3 appears in Section 13.4.

For $G_0 \in \mathfrak{S}_{n,d}$, *exact* random k-colouring can be implemented efficiently. In what follows we describe an efficient process that can colour randomly any graph in $\mathfrak{S}_{n,d}$.

Random Colouring in $\mathfrak{S}_{n,d}$

Input: $G \in \mathfrak{S}_{n,d}$, k.

`set` $\mathcal{C}$ to be the set of all cycles in G

`for` each isolated vertex $v \in V(G)$ `do` /*Colouring isolated vertices*/

  `set` $\tau(v)$ a colour chosen uniformly random from $[k]$

`for` each $C = (w_0, \ldots w_l) \in \mathcal{C}$ `do` /*Colouring isolated cycles*/

  `set` $\tau(w_0)$ a color chosen uniformly random from $[k]$

  `for` $i = 1, \ldots, l$ `do`

    `set` μ_{w_i} the Gibbs marginal of w_i, conditional $\tau(w_0), \ldots, \tau(w_{i-1})$

    `compute` μ_{w_i} using Dynamic Programming

    `set` $\tau(w_i)$ according to μ_{w_i}

Output: τ

The most interesting part of the above algorithm is the one for random colouring of the cycles. For each cycle $C \in \mathcal{C}$, the algorithm first assigns a random colour on the vertex w_0. Once w_0 is assigned a colour, then we eliminate the cycle structure of C and now we deal with a tree of maximum degree 2. This allows to compute the marginal μ_{w_i}, for each vertex $w_i \in C$, by using *Dynamic Programming* (DP).

Remark 8 *The use of DP for computing Gibbs marginals on the trees is well known to be exact, e.g. see [19] for an excellent survey on the subject.*

Remark 9 *The recursive distributional equations that DP uses in this setting are more or less standard. Example of such equations appear in the proof of Lemma 6, in Section 11.1.*

Once we get an exact random colouring of G_0 by using the above algorithm, `Random Colouring Algorithm` colours the remaining graphs $G_1, \ldots, G_r$ by using `Update`, as we described in Section 5.

Let $\mathfrak{X}_{n,d}$ contain every graph G on n vertices such that the following holds:

1. getting a sequence of subgraphs $G_0, \ldots, G_r$, as described in Section 6, it holds that $G_0 \in \mathfrak{S}_{n,d}$
2. $|E(G)| \leq (1 + n^{-1/3})dn/2$.

[4] single isolated vertex

[5] the cycles do not share edges nor vertices

Note that for some G we have that $G_0 \in \mathfrak{S}_{n,d}$ regardless of the order we remove the edges for creating the sequence $G_0, \ldots, G_r$. That is, whether $\boldsymbol{G} \in \mathfrak{X}_{n,d}$, or not, depends only on the graph $\boldsymbol{G}$.

If the input graph $\boldsymbol{G}$ does not belong into $\mathfrak{X}_{n,d}$, then the `Random Colouring Algorithm` abandons. It turns out that this typically does not happen. In particular, we have following corollary.

Corollary 1 *For sufficiently large $d > 0$, it holds that* $\Pr[\boldsymbol{G} \in \mathfrak{X}_{n,d}] \geq 1 - 2n^{-2/3}$.

Proof: Lemma 3, states that for the sequence $G_0, \ldots, G_r$ generated from $\boldsymbol{G}$ as described in Section 6 it holds that $\Pr[G_0 \in \mathfrak{S}_{n,d}] \geq 1 - n^{-2/3}$. Using Chernoff's bounds, e.g. [13], we also get

$$\Pr\left[|E(\boldsymbol{G})| \geq (1 + n^{-1/3})dn/2\right] \leq \exp\left(-n^{1/4}\right).$$

A simple union bound, yields that indeed $\Pr[\boldsymbol{G} \in \mathfrak{X}_{n,d}] \geq 1 - 2n^{-2/3}$. $\Diamond$

In the following two sections we prove Theorems 1 and 2.

6.1 Proof of Theorem 1

For proving Theorem 1 we need to use the following result, whose proof appears in Section 7.

Theorem 8 *Let ϵ, d, k be as in the statement of Theorem 1. Consider the sequence $G_0, \ldots, G_r$ generated from $\boldsymbol{G}$ as described in Section 6. For any $i \in \{0, \ldots, r-1\}$ it holds that*

$$\mathbb{E}[\alpha_i] \leq 50\epsilon^{-1}k(4+\epsilon)n^{-\left(1+\frac{\epsilon}{36(1+\epsilon/4)\log d}\right)}.$$

Proof of Theorem 1: In light of Corollary 1, it suffices to show that (1) holds with sufficiently large probability over the instances $\boldsymbol{G}$, conditional that $\boldsymbol{G} \in \mathfrak{X}_{n,d}$.

Let $\mathcal{A}$ be the event $\boldsymbol{G} \in \mathfrak{X}_{n,d}$. First we argue about $\mathbb{E}\left[||\mu - \hat{\mu}|| \mid \mathcal{A}\right]$, i.e. the expectation is w.r.t. the instances $\boldsymbol{G}$. Using Theorem 7 and Theorem 8 we have that

$$\mathbb{E}\left[||\mu - \hat{\mu}|| \mid \mathcal{A}\right] \leq \mathbb{E}\left[\sum_{i=0}^{r-1} \alpha_i \mid \mathcal{A}\right],$$

where the expectation is taken over the instances $\boldsymbol{G}$. Noting that $\alpha_i \in [0,1]$, we get

$$\mathbb{E}\left[||\mu - \mu'|| \mid \mathcal{A}\right] \leq \sum_{i=0}^{(1+n^{-1/3})dn/2} \mathbb{E}[\alpha_i \mid \mathcal{A}], \tag{4}$$

where the above follows by observing that $\mathcal{A}$ implies that $r \leq (1 + n^{-1/3})dn/2$.

On the other hand for the quantities $\mathbb{E}[\alpha_i \mid \mathcal{A}]$ we work as follows:

$$\begin{aligned} \mathbb{E}[\alpha_i \mid \mathcal{A}] &\leq (\Pr[\mathcal{A}])^{-1} \cdot \mathbb{E}[\alpha_i] && \text{[since } \alpha_i \geq 0\text{]} \\ &\leq 100\epsilon^{-1}k(4+\epsilon)n^{-\left(1+\frac{\epsilon}{36(1+\epsilon/4)\log d}\right)}, && (5) \end{aligned}$$

in the final inequality we used Theorem 8 and Corollary 1. Plugging (5) into (4), we get that

$$\mathbb{E}\left[||\mu - \hat{\mu}|| \mid \mathcal{A}\right] \leq C \cdot n^{-\frac{\epsilon}{36(1+\epsilon/4)\log d}},$$

for fixed $C > 0$. The theorem follows by applying Markov's inequality. $\Diamond$

6.2 Proof of Theorem 2

First, we are going to show that, on input $G \in \mathfrak{X}_{n,d}$, `Random Colouring Algorithm` has time complexity $O(n^2)$. Then, the theorem will follow by using Corollary 1.

We start by considering the time complexity of the algorithm on input $G \in \mathfrak{X}_{n,d}$. First the algorithm constructs $G_0, \ldots, G_r$. For this, it needs to distinguish which edges in $E(G)$ do not belong to a short cycle. This can be done by exploring the structure of the $(\log_d n)/9$-neighbourhood around each edge of G by using *Breadth First Search* (BFS). The search around each edge requires $O(n)$ steps, since $|E(G)| = O(n)$. The exploration is repeated for each edge in $E(G)$. Thus, the algorithm requires $O(n^2)$ steps to find the short cycles. This implies that $G_0, \ldots, G_r$ can be constructed in $O(n^2)$ steps.

Since the $|E(G_i)| = O(n)$, for every $i = 0, \ldots, r$, Theorem 4 implies that the number of steps required for each `Update` call is $O(n)$. Consequently, we need $O(n^2)$ steps for all the calls of `Update`, since $r \leq |E(G)| = O(n)$.

It remains to consider the time complexity of colouring randomly G_0. The algorithms uses `Random Colouring in` $\mathfrak{S}_{\mathtt{n,d}}$ (Section 6) to colour randomly G_0. Due to our assumptions it holds that $G_0 \in \mathfrak{S}_{n,d}$. Let $\mathcal{C}$ be the set of cycles in G_0. Note that all the cycles in $\mathcal{C}$ are simple and isolated from each other. Also, all the vertices in G_0 which are not in a cycle are isolated.

We consider the time complexity of colouring the cycles in $\mathcal{C}$. For each $C = (w_0, \ldots, w_{|C|}) \in \mathcal{C}$, first, the problem is reduced to computing Gibbs marginals on a tree of maximum degree 2. This is done by assigning w_0 a uniformly random colour from $[k]$. Then, the algorithm colours iteratively the vertices in C. At iteration i, the colouring of the vertices $w_1, \ldots, w_{i-1}$ is already known and the algorithm colours w_i as follows: It computes the marginal μ_{w_i}, conditional the colour assignment of the vertices $w_0, \ldots, w_{i-1}$, by using Dynamic Programming. Then it assigns a colour to w_i according to μ_{w_i}.

Given the distribution of the children of w_i w.r.t. the subtree that hangs from them, the Dynamic Program requires $O(k^2)$ arithmetic operations to compute μ_{w_i}. This means that the algorithm requires $O(k^2|C|)$ operations for computing μ_{w_i}. It is clear that each cycle C requires at most $O(k^2|C|^2)$ steps to be coloured randomly.

Consequently, the algorithm requires $O(k^2 n \log^2 n)$ number of steps to colour randomly all the cycles in $\mathcal{C}$, since $|C| = O(\log n)$ and $|\mathcal{C}| = O(n)$. Additionally, the algorithm requires $O(n)$ steps to colour randomly all the $O(n)$ many isolated vertices.

Concluding, the time complexity of `Random Colouring in` $\mathfrak{S}_{\mathtt{n,d}}$, for fixed k is $O(n \log^2 n)$. This implies that `Random Colouring Algorithm`, on input $G \in \mathfrak{E}_{n,d}$, has time complexity $O(n^2)$.

The theorem follows.

7 Proof of Theorem 8

Let $\Lambda_{n,k}$ denote the set of all the 4-tuples (G, v, u, σ) such that G is a k colourable graph on n vertices, $v, u \in V(G)$ and σ is a k-colouring of G. For $(G, v, u, \sigma) \in \Lambda_{n,k}$ and $q \in [k] \backslash \{\sigma_v\}$, consider the disagreement graph $\boldsymbol{Q} = \boldsymbol{Q}(G, v, \sigma, q)$ and let the event $\mathcal{Q}_{\sigma_v, q} =$ "$u \in \boldsymbol{Q}$".

For $c_1, c_2 \in [k]$ and an integer $i \geq 0$ we let the distribution $\mathcal{P}^i_{c_1,c_2}$ over $(G, v, u, Z) \in \Lambda_{n,k}$ be induced by the following experiment: Take an instance $\boldsymbol{G}$ and construct the sequence $G_0, \ldots, G_r$ as described in Section 6. Then,

1. G is equal to G_i
2. v and u are equal to v_i and u_i, respectively
3. Z is distributed uniformly at random in $\Omega_G(c_1, c_2)$

Remark 10 *In $G_0, \ldots, G_r$, the number of terms in the sequence is a random variable. In the definition of $\mathcal{P}^i_{c_1,c_2}$ if $i > r$ we follow the convention that G is the empty graph with probability* 1.

Also, denote by $\mathcal{P}^i_{*,c_2}$ the distribution when $Z(v)$ is not fixed, i.e. Z is a random k-colouring of G, conditional that $Z(u) = c_2$. In the same manner, denote by $\mathcal{P}^i_{c_1,*}$, the distribution when $Z(u)$ is not fixed. Finally, we define $\mathcal{P}^i_{*,*}$ when there is no restriction on the colouring of both v, u.

For proving Theorem 8 we need the following two results.

Proposition 1 *Let ϵ, d and k be as in the statement of Theorem 8. Let $c, q \in [k]$ be such that $c \neq q$. For any $i \geq 0$, it holds that*

$$\mathcal{P}^i_{c,*}[\mathcal{Q}_{c,q}] \leq 10\epsilon^{-1}(4+\epsilon)n^{-\left(1+\frac{\epsilon}{36(1+\epsilon/4)\log d}\right)}.$$

The proof of Proposition 1 appears in Section 8.

Lemma 4 *Let ϵ, d, k be as in the statement of Theorem 8. For any $c \in [k]$ and any $i \geq 0$ it holds that*

$$||\mathcal{P}^i_{c,*}(\cdot) - \mathcal{P}^i_{*,*}(\cdot)||_{\{u_i\}} \leq n^{-1}.$$

The proof of Lemma 4 appears in Section 7.1.

Proof of Theorem 8: It is elementary to verify that

$$\mathbb{E}[\alpha_i] \leq \max_{c,q\in[k]:c\neq q} \left\{\mathcal{P}^i_{c,c}[\mathcal{Q}_{c,q}] + \mathcal{P}^i_{q,c}[\mathcal{Q}_{q,c}]\right\}. \tag{6}$$

The theorem follows by bounding appropriately the probability terms in the r.h.s. of (6).

Given $(G, v, u, \sigma) \in \Lambda_{n,k}$, we let the events $E :=$"$\sigma(v) = \sigma(u)$" and $A_{c_1} :=$"$\sigma(u) = c_1$", for every $c_1 \in [k]$. Since it holds that $\mathcal{P}^i_{c,*}[\mathcal{Q}_{c,q}] \geq \mathcal{P}^i_{c,*}[\mathcal{Q}_{c,q}|E] \cdot \mathcal{P}^i_{c,*}[E]$ and $\mathcal{P}^i_{c,*}[\cdot|E] = \mathcal{P}^i_{c,c}[\cdot]$, we get that

$$\mathcal{P}^i_{c,c}[\mathcal{Q}_{c,q}] \leq \frac{1}{\mathcal{P}^i_{c,*}[E]}\mathcal{P}^i_{c,*}[\mathcal{Q}_{c,q}]. \tag{7}$$

Noting that $\mathcal{P}^i_{c,*}[E] = \mathcal{P}^i_{c,*}[A_c]$ and $\mathcal{P}^i_{*,*}[A_c] = k^{-1}$, from Lemma 4 we get that

$$\left|\mathcal{P}^i_{c,*}[E] - k^{-1}\right| \leq n^{-1}. \tag{8}$$

Using (8) and (7) we get that

$$\mathcal{P}^i_{c,c}[\mathcal{Q}_{c,q}] \leq 2k \cdot \mathcal{P}^i_{c,*}[\mathcal{Q}_{c,q}] \leq 20\epsilon^{-1}k(4+\epsilon)n^{-\left(1+\frac{\epsilon}{36(1+\epsilon/4)\log d}\right)}, \tag{9}$$

where the last inequality follows from Proposition 1. Applying the same arguments we, also, get that

$$\mathcal{P}^i_{q,c}[\mathcal{Q}_{q,c}] \leq 20\epsilon^{-1}k(4+\epsilon)n^{-\left(1+\frac{\epsilon}{36(1+\epsilon/4)\log d}\right)}. \tag{10}$$

The bounds in (9) and (10) hold for any $c, q \in [k]$, different with each other. The theorem follows by plugging (9) and (10) into (6). ♢

7.1 Proof of Lemma 4

Let $(G, v, u, X), (G, v, u, Z) \in \Lambda_{n,k}$, for some fixed G. Let X, Z be two coupled random colourings of G. In particular for X, Z we have the following: Assuming that $X(v) = c$, we choose q u.a.r. among $[k]$ and we set $Z(v) = q$. Depending on whether $c = q$ or not the coupling does the following choices.

Case "$c = q$": Couple Z and X identically, i.e. $X = Z$

Case "$c \neq q$": Set $Z = \texttt{Switching}(G, v, X, q)$,

where Switching is from Section 4. Claim 1 establishes that Z follows the appropriate distribution.

Claim 1 $\texttt{Switching}(G, v, X, q)$ *is a random colouring of* G *conditional on that* v *is coloured* q.

Proof: It suffices to show that the sets $\Omega_c = \cup_{c' \in [k]} \Omega_i(c, c')$ and $\Omega_q = \cup_{c' \in [k]} \Omega_i(q, c')$ admit the bijection $\texttt{Switching}(G, v, \cdot, q) : \Omega_c \to \Omega_q$.

First, note that Lemma 1 implies that if $\tau = \texttt{Switching}(G, v, \sigma, q)$, then $\tau \in \Omega_{G,k}$. Also, it is direct that $\tau \in \Omega_q$. Second, we need to show that the mapping $\texttt{Switching}(G, v, \cdot, q) : \Omega_c \to \Omega_q$ is *surjective*, i.e. for any $\sigma \in \Omega_q$ there is a $\sigma' \in \Omega_c$ such that $\sigma = \texttt{Switching}(G, v, \sigma', q)$. Clearly, such σ' exists. In particular, it holds that $\sigma' = \texttt{Switching}(G, v, \sigma, c)$. The last observation also implies that the mapping is *one-to-one*. Since $\texttt{Switching}(G, v, \cdot, c)$ is surjective and one-to-one it is a bijection. The claim follows. $\Diamond$

For the case where $q \neq c$, consider the disagreement graph $\boldsymbol{Q} = \boldsymbol{Q}(G, v, X, q)$. We remind the reader that the event $\mathcal{Q}_{c,q}$:="$u \in \boldsymbol{Q}$". Due to the way we construct Z we have that the event $\mathcal{Q}_{c,q}$ holds if and only if $X(u) \neq Z(u)$ holds. That is,

$$\Pr[X(u) \neq Z(u)] \leq \Pr[\mathcal{Q}_{c,q}]. \tag{11}$$

Note that the probability terms above hold for any k-colourable graph G.

For our purpose, we need to consider $(G, v, u, X), (G, v, u, Z)$ distributed as in $\mathcal{P}^i_{c,*}$ and $\mathcal{P}^i_{q,*}$ respectively, for $q \neq c$. For such 4-tuples, (11) implies that

$$\Pr[X(u) \neq Z(u)] \leq \mathcal{P}^i_{c,*}[Q_{c,q}].$$

Note that the above is derived by taking averages w.r.t. the graph instance G_i in the sequence $G_0, \ldots, G_r$ where (v, u) correspond to (v_i, u_i). The lemma follows by noting that

$$||\mathcal{P}^i_{c,*}(\cdot) - \mathcal{P}^i_{*,*}(\cdot)||_{\{u\}} \leq \mathcal{P}^i_{c,*}[Q_{c,q}],$$

while from Proposition 1 we have that $\mathcal{P}^i_{c,*}[Q_{c,q}] \leq n^{-1}$.

8 Proof of Proposition 1

Let (G, v, u, X) be distributed as in $\mathcal{P}^i_{c,*}$. Every path P in G which start from v and $\forall w \in P$ we have $X(w) \in \{c, q\}$ is called *path of disagreement*. It holds that

$$\mathcal{P}^i_{c,*}[\mathcal{Q}_{c,q}] \leq \mathcal{P}^i_{c,*}[B] + \mathcal{P}^i_{c,*}[C],$$

where the events B and C are as follows: B :="v and u are connected through a path of disagreement of length at most $\log^2 n$". C :="v and u are connected through a path of length greater than $\log^2 n$'.

Let, also, the event C' := "there is a path of disagreement starting from v and has length greater than $\log^2 n$". Note that the event C' does not specify the end vertex of the path of disagreement. It is immediate that $\mathcal{P}^i_{c,*}[C'] \geq \mathcal{P}^i_{c,*}[C]$, since, the event C is included in the event C'. Thus, it holds that

$$\mathcal{P}^i_{c,*}[\mathcal{Q}_{c,q}] \leq \mathcal{P}^i_{c,*}[B] + \mathcal{P}^i_{c,*}[C'].$$

The proposition will follow by bounding appropriately the probabilities $\mathcal{P}^i_{c,*}[B]$ and $\mathcal{P}^i_{c,*}[C']$.

For every vertex w, we let $\Gamma_w(l)$ denote the number of paths of disagreement of length l that connect v and w. From Markov's inequality we get that

$$\mathcal{P}^i_{c,*}[B] \leq \mathbb{E}_{\mathcal{P}^i_{c,*}} \left[\sum_{l \leq \log^2 n} \Gamma_u(l) \right], \tag{12}$$

where $\mathbb{E}_{\mathcal{P}^i_{c,*}}[\cdot]$ is the expectation w.r.t. $\mathcal{P}_{i_{c,*}}$. For bounding $\mathcal{P}^i_{c,*}[C']$ we use the following inequality

$$\mathcal{P}^i_{c,*}[C'] \leq \mathbb{E}_{\mathcal{P}^i_{c,*}}\left[\sum_w \Gamma_w(\log^2 n)\right], \tag{13}$$

where the summation on the r.h.s. of the inequality, above, runs over all the vertices of the graph.

So as to compute the expectation both in (12) and (13) we use Theorem 3. However, we note that the pair of vertices v, u we consider is not a uniformly random one. Since we consider the probability distribution $\mathcal{P}^i_{c,*}$, the pair v, u is distributed uniformly at random among the pair of vertices which are at distance greater than $(\log_d n)/9$ in $\boldsymbol{G}$.

Letting p be the probability that a randomly chosen edge from $\boldsymbol{G}$ does not belong to a cycle of length less than $(\log_d n)/9$. Using Theorem 3 we get that

$$\mathbb{E}_{\mathcal{P}^i_{c,*}}\left[\sum_{l\leq \log^2 n} \Gamma_u(l)\right] \leq 2p^{-1}\sum_{l\geq l_0}^{\log^2 n} n^{l-1}\left((1+\epsilon/4)n\right)^{-l}, \qquad \text{for } l_0 = (\log_d n)/9 + 1. \tag{14}$$

Let us explain how do we get the above inequality from Theorem 3. If the vertices v, u were not conditioned to be at distance greater than $(\log_d n)/9$, then the expected number of paths of disagreement of length l between them is equal to the number of possible paths of length l times the probability each of these paths is a path of disagreement. Clearly the number of the possible paths is at most n^{l-1}, i.e. we have fixed the first and the last vertex of the paths. From Theorem 3 we have that the probability of each of these paths to be disagreeing is $2\left((1+\epsilon/4)n\right)^{-l}$. We divide by p due to conditioning that the vertices v, u are not entirely random, since we have conditioned that their distance is larger than $(\log_d n)/9$.

It is direct to show that it holds that $p \geq 1 - n^{-9/10}$. Then, we have that

$$\mathbb{E}_{\mathcal{P}^i_{c,*}}\left[\sum_{l\leq \log^2 n} \Gamma_u(l)\right] \leq 4\epsilon^{-1}(4+\epsilon)n^{-1-\frac{\epsilon}{36(1+\epsilon/4)\log d}}. \tag{15}$$

Working in the same manner for (13) we get that

$$\begin{aligned}\mathbb{E}_{\mathcal{P}^i_{c,*}}\left[\sum_w \Gamma_w(\log^2 n)\right] &\leq 2p^{-1}\left(1+\epsilon/4\right)^{-\log^2 n} \\ &\leq 2p^{-1}n^{-((\log n)\cdot\log(1+\epsilon/4))} \leq n^{-\sqrt{\log n}},\end{aligned} \tag{16}$$

where the last inequality holds for large n and noting that $p > 1/2$. Observe that in the second case the number of paths of length l that emanate from v is at most n^l, as we do not fix the last vertex of the path.

Using (15) and (12) we bound appropriately $\mathcal{P}^i_{c,*}[B]$. Using (16) and (13) we bound appropriately $\mathcal{P}^i_{c,*}[C']$. The proposition follows.

9 Proof of Theorem 3

For the sake of brevity we denote with P not only the permutation of the vertices $w_0, \ldots, w_\ell$ but the corresponding path in $\boldsymbol{H}$, if such path exists. The probability term in (3) is w.r.t. both the randomness of the graph $\boldsymbol{H}$ and the random k-colourings of $\boldsymbol{H}$. That is, for $\mathbf{I}_{\{P\}} = 1$, first we need to have that the vertices in the permutation P form path in $\boldsymbol{H}$. Then, given that $\boldsymbol{H}$ contains the path P, we need to bound the probability that this path is 2-coloured in a random k-colouring of $\boldsymbol{H}$. Clearly, the challenging part is the second one. We denote $\boldsymbol{H}_P$ the graph $\boldsymbol{H}$ conditional that the path P appears in the graph.

Our approach is as follows: Given $\boldsymbol{H}_P$, first we specify an appropriate subgraph of $\boldsymbol{H}_P$ which includes the path P. We call this subgraph $\boldsymbol{N}$. Also, we specify a set $\boldsymbol{B} \subset V(\boldsymbol{N})$ such that $\boldsymbol{B}$ separates $V(\boldsymbol{N})\backslash\boldsymbol{B}$ from the rest of the graph $\boldsymbol{H}_P$. We set an appropriate (worst case) boundary condition $\sigma_{\boldsymbol{B}} \in [k]^{\boldsymbol{B}}$ on $\boldsymbol{B}$. Let μ_N^σ, be the Gibbs distribution of the k-colourings of $\boldsymbol{N}$, conditional that $\boldsymbol{B}$ is coloured σ_B. The choice of σ is such that under μ_N^σ the probability of P to be 2-coloured with c, q is lower bounded by the corresponding probability under μ_H, the Gibbs distribution of the k-colourings of $\boldsymbol{H}_P$.

Let us describe how do we get $\boldsymbol{N}$ and $\boldsymbol{B} \subset V(\boldsymbol{N})$. For this, we consider an integer parameter $h = h(\epsilon) > 0$, which we assume that is sufficiently large it depends on ϵ and it is independent of d.

Path Neighbourhood Revealing. Consider the graph $\boldsymbol{H}_P$. For each $w_i \in P$ we define the sets $L_{i,s} \subseteq V(\boldsymbol{H}_P)$, for $s = 0, \ldots, h$, as follows: $L_{i,0} = \{w_i\}$. We get $L_{i,s}$ by working inductively, i.e. we use $L_{i,s-1}$. Let $\mathcal{R}_{i,s} \subset V(\boldsymbol{G})$ contain all the vertices but those which belong to P and those which belong in $\bigcup_{j<i}\bigcup_{j'\leq h} L_{j,j'}$ and $\bigcup_{j'<s} L_{i,j'}$. Consider an (arbitrary) ordering of the vertices in $\mathcal{R}_{i,s}$. For each vertex $u \in L_{i,s-1}$ we examine its adjacency with the vertices in $\mathcal{R}_{i,s}$ in the predefined order. We stop revealing the neighborhood of u in $\mathcal{R}_{i,s}$ once we either have revealed $(1+\epsilon/3)d+1$ many neighbours, or if we have checked all the possible adjacencies of u with $\mathcal{R}_{i,s}$. Whichever happens first[6]. Then $L_{i,s}$ contains all the vertices in $\mathcal{R}_{i,s}$ which have been revealed to have a neighbour in $L_{i,s-1}$.

For $i = 0, \ldots, \ell$, let $N_{i,h}$ be the induced subgraph of $\boldsymbol{H}_P$ with vertex set $\bigcup_{s=0}^{h} L_{i,s}$. Note that the size of $N_{i,h}$ depends only on ϵ, d, h, i.e. it is independent of n. In particular, it holds that

$$|V(N_{i,h})| \leq N_0 = \frac{[(1+\epsilon/3)d+1]^{h+1}-1}{(1+\epsilon/3)d}. \tag{17}$$

We call $N_{i,h}$, `Fail` if at least one of the following happens:

- The maximum degree in $N_{i,h}$ is at least $(1+\epsilon/3)d+1$
- The graph $N_{i,h}$ is not a tree
- There is an integer $j \neq i$ such that some vertex $w'' \in N_{j,h}$ is adjacent to some vertex $w' \in \mathcal{N}_{i,h}$ and the edge $\{w', w''\}$ does not belong to the path P.

Lemma 5 *Let ϵ, d be as in Theorem 3. Consider a sufficiently large fixed integer $h = h(\epsilon) > 0$, independent of d. Let F be the number of vertices $w_i \in P$ such that $N_{i,h}$ is* `Fail`*, for $i = 1, \ldots, \ell$. For any $s = 1, \ldots, \ell$, it holds that*

$$\Pr[F = s] \leq (1+n^{-1/3})\binom{\ell}{s}\exp\left[-\epsilon^2 ds/35\right].$$

In the lemma, above, F does not consider $N_{0,h}$. The proof of Lemma 5 appears in Section 9.1.

The graph $\boldsymbol{N}$ we are looking for is a subgraph of $\bigcup_{i=0}^{\ell} N_{i,h}$. For specifying $\boldsymbol{N}$ perhaps it is more natural to start with the set $\boldsymbol{B}$ which separates $\boldsymbol{N}$ from the rest of $\boldsymbol{H}_P$. Each time, we decide on $\boldsymbol{B} \cap V(N_{i,h})$ by examining each $N_{i,h}$, separately. If $N_{i,h}$ is `Fail`, then $\boldsymbol{B} \cap V(N_{i,h}) = \{w_i\}$, i.e the vertex in the path P. On the other hand, if $N_{i,h}$ is not `Fail`, then $\boldsymbol{B} \cap V(N_{i,h}) = L_{i,h}$, i.e. all the vertices in $N_{i,h}$ that are at distance h from w_i.

In Figure 7, we see one example of a possible outcome of the exploration we describe above. The lined vertices are exactly those which belong to the boundary set $\boldsymbol{B}$. If some vertex w_i on the path is lined, this means that $N_{i,h}$ is `Fail`. The vertices of the path which are not lined correspond to the roots of a "low degree" tree of height at most h.

[6] Clearly, as the process goes, the number of neighbours of u in $\mathcal{R}_{i,s}$ is at most $(1+\epsilon/3)d$+1.

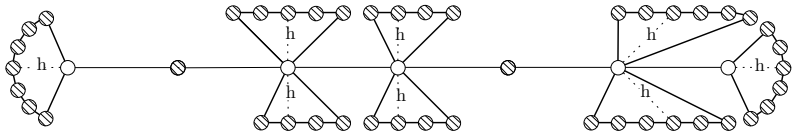

Figure 7: The lined vertices belong to $\boldsymbol{B}$.

Let $S \subseteq \{0, \ldots, \ell\}$ contain each i such that $N_{i,h}$ is not `Fail`. Also, let $V_A = \bigcup_{i \in S} V(N_{i,h-1})$ [7]. It is not hard to see that the vertex set $\boldsymbol{B}$ is a cut-set that separates V_A from the rest of the vertices in $V(\boldsymbol{H}_P)$. The graph $\boldsymbol{N}$ is the induced subgraph of $\boldsymbol{H}_P$ with vertex set $V_A \cup \boldsymbol{B}$.

Remark 11 *Since $\boldsymbol{H}_P$ is random, the subgraph $\boldsymbol{N}$ is random.*

Consider the graph $\boldsymbol{H}_P$ and the corresponding Gibbs distribution μ_H. The distribution μ_H specifies a convex combination of boundary conditions on $\boldsymbol{B}$. Using these boundary conditions we could estimate the probability that P is coloured only with c, q, exactly. However, estimating this convex combination of boundaries is a formidable task to accomplish. We get an upper bound of this probability by considering a worst boundary condition on the vertex set $\boldsymbol{B}$. The condition is worst in the sense that it maximizes the probability of interest. That is, instead of μ_H, we consider the distribution μ_N^σ which is much easier to handle. Under μ_N^σ the probability that P is coloured with c, q is at least as big as under μ_H.

In the following results, we let $\mathcal{T}_{d,\epsilon,h}$ be the set of labeled, rooted, trees of maximum degree $(1 + \epsilon/3)d$ and height h.

Proposition 2 *Let ϵ, d, k be as in Theorem 3. Consider a sufficiently large fixed integer $h = h(\epsilon) > 0$, independent of d. Consider $\boldsymbol{H}_P$ and let $\boldsymbol{N}, \boldsymbol{B}$ be as defined above. For each $w_j \in P$ such that $w_j \notin \boldsymbol{B}$ the following is true:*

Let Γ be the neighbours of w_j in the path P and let $\boldsymbol{B}^+ = \boldsymbol{B} \cup \Gamma$. There exists a function $f_\epsilon : \mathbb{N} \to \mathbb{R}^+$, such that $f(h) \to 0$ as $h \to \infty$, while for any $\sigma \in \Omega_{\boldsymbol{N},k}$ and any $c \in [k]$ it holds that

$$\max_{N_{j,h} \in \mathcal{T}_{d,\epsilon,h}} |\Pr[X(v_j) = c \mid N_{j,h}, X_{\boldsymbol{B}^+} = \sigma_{\boldsymbol{B}^+}] - \Pr[X(v_j) = c \mid N_{j,h}, X_\Gamma = \sigma_\Gamma]| \leq k^{-1} f_\epsilon(h),$$

where X is a random k-colouring of $\boldsymbol{N}$.

Note that the above is a spatial mixing result. It implies that for any $N_{j,h}$ which is not `Fail` the boundary we set at distance h from w_j, essentially, has no effect on the distribution of the k-colouring of w_i. The proof of Proposition 2 appears in Section 10.

For every $w_j \in P$ such that $w_j \in \boldsymbol{B}$, the worst case boundary condition sets the vertex to its appropriate colour, i.e. if j is even then the colour is c, otherwise the colour is q. Proposition 2 implies that, whatever is the boundary condition at $\boldsymbol{B}$, if $w_j \notin \boldsymbol{B}$, its probability of getting colour q or c, depending on the parity of j, is approximately $1/k$.

Proof of Theorem 3: Let E_P be the event that $\boldsymbol{H}$ contains the path P. It holds that

$$\Pr\left[\mathbf{I}_{\{P\}} = 1\right] \leq (d/n)^\ell \cdot \Pr\left[\mathbf{I}_{\{P\}} = 1 \mid E_P\right].$$

Consider $\boldsymbol{H}_P$ and let X be a random k-colouring conditional on that $X(w_0) = c$. For i even, we call $w_i \in P$ disagreeing if $X(w_i) = c$. For i odd number, we call $w_i \in P$ disagreeing if $X(w_i) = q$.

Let the event D_i that "w_i is disagreeing". Clearly it holds that

$$\Pr\left[\mathbf{I}_{\{P\}} = 1\right] \leq (d/n)^\ell \Pr\left[\cap_{i=1}^{\ell} D_i \mid E_P\right]. \tag{18}$$

[7] $N_{i,h-1}$ is defined in the natural way

Let the events A_i, B_i, C_i be defined as follows: $A_i =$ "$N_{i,h}$ is `Fail`". $B_i =$ "$N_{i,h}$ is not `Fail` and w_i is disagreeing". Also let $C_i = A_i \cup B_i$.

Claim 2 *It holds that*

$$\Pr\left[\cap_{i=1}^{\ell} D_i \mid E_P\right] \leq \Pr\left[\cap_{i=1}^{\ell} C_i \mid E_P\right].$$

Proof: In the setting of the proof of Theorem 3, assume that we have revealed the underlying graph $\boldsymbol{H}_P$. It suffices to show that

$$\Pr\left[\cap_{i=1}^{\ell} D_i \;\middle|\; \boldsymbol{H}_P\right] \leq \Pr\left[\cap_{i=1}^{\ell} C_i \;\middle|\; \boldsymbol{H}_P\right]. \tag{19}$$

Observe that the probability terms are only w.r.t. the random colouring of $\boldsymbol{H}_P$.

Let W be the set of vertices $q_i \in P$ such that $N_{i,h}$ is not `Fail`. Also, let $W' \subseteq \boldsymbol{B}$ be the set of vertices $w_i \in P$ for which $N_{i,h}$ is `Fail`. The events $\cap_{w_i \in W} C_i$ and $\cap_{w_i \in W} D_i$ are identical, since both occur if the vertices in W are disagreeing. Thus it holds that $\Pr\left[\cap_{w_i \in W} D_i \mid \boldsymbol{H}_P\right] = \Pr\left[\cap_{w_i \in W} C_i \mid \boldsymbol{H}_P\right]$.

Furthermore, we note that $\Pr\left[\cap_{w_i \in W'} C_i \mid \boldsymbol{H}_P, \cap_{w_i \in W} C_i\right] = 1$. On the other hand, it holds that $\Pr\left[\cap_{w_i \in W'} D_i \mid \boldsymbol{H}_P, \cap_{w_i \in W} D_i\right] \leq 1$. These imply that (19) is true. The claim follows. ◊

Using Claim 2 and (18), it suffices to bound appropriately $\Pr\left[\cap_{i=1}^{\ell} C_i \mid E_P\right]$.

Consider $\boldsymbol{H}_P$ and let $\mathcal{F}_i(C)$ be the σ-algebra generated by the events C_j, for every $j \neq i$. Proposition 2 implies that

$$\rho = \Pr[B_i \mid \mathcal{F}_i(C), E_P, N_{i,h} \text{ is not } \texttt{Fail}] \leq (k-2)^{-1} + f_\epsilon(h)/k. \tag{20}$$

for any $i = 0, \ldots, \ell$. Letting F be the number of vertices $w_i \in P$ such that $N_{i,h}$ is `Fail`, for $i = 1, \ldots, \ell$, we have that

$$\begin{aligned}
\Pr\left[\cap_{i=1}^{\ell} C_i \mid E_P\right] &= \sum_{s=0}^{\ell} \Pr\left[\cap_{i=1}^{\ell} C_i \mid E_P, F = s\right] \Pr\left[F = s \mid E_P\right] \\
&\leq \sum_{s=0}^{\ell} \rho^{\ell - s} \Pr\left[F = s \mid E_P\right] && \text{[from (20)]} \\
&\leq (1 + n^{-1/3}) \sum_{s=0}^{\ell} \binom{\ell}{s} \rho^{\ell - s} \exp(-\epsilon^2 d s/35) && \text{[from Lemma 5]} \\
&\leq 2\left[\rho + \exp(-\epsilon^2 d/35)\right]^{\ell}. && (21)
\end{aligned}$$

Using the fact that $k \geq (1+\epsilon)d$, for sufficiently large h, d, (21) implies that

$$\Pr\left[\cap_{i=1}^{\ell} C_i \mid E_P\right] \leq 2((1+\epsilon/4)d)^{-\ell}. \tag{22}$$

The theorem follows from (22), (18) and Claim 2. ◊

9.1 Proof of Lemma 5

For proving the lemma we use the following tail bound, [13], Corollary 2.3. Let W be distributed as in $\mathcal{B}(n, d/n)$, i.e. binomial distribution with parameters n and d/n. For any fixed $\alpha > 0$ and sufficiently large d, it holds that

$$\Pr[W \geq (1+\alpha)d] \leq \exp\left(-\alpha^2 d/3\right). \tag{23}$$

For $i, j = 0, \ldots, \ell$ consider the following events: Let $A_i :=$" $N_{i,h}$ has maximum degree greater than $(1+\epsilon/3)d$". Also, let $B_i :=$"$N_{i,h}$ is not a tree". For any two i, j such that $i \neq j$, we let $E_{i,j} :=$"there is an edge, not in P, which connects some vertex in $N_{i,h}$ and some vertex in $N_{j,h}$".

Given some $i \in \{0,\ldots,\ell\}$ and any $S \subset \{0,\ldots,\ell\}$ such that $i \notin S$, let $\mathcal{F}_S$ be the σ-algebra generated be the events A_j, B_j for $j \in S$. Given, $\mathcal{F}_S$, for every vertex $w \in L_{i,t-1}$ has a number of neighbours in $\mathcal{R}_{i,t}$ which is dominated by $\mathcal{B}(n, d/n)$, for $t = 1,\ldots,h$. Then, (23) implies that the probability for w to have at least $(1+\epsilon/3)d$ neighbours in $\mathcal{R}_{i,t}$ is at most $\exp\left(-\epsilon^2 d/27\right)$.

The event A_i occurs if there exists $t \in [h]$ and $w \in L_{i,t-1}$ whose number of neighbour in $\mathcal{R}_{i,t}$ is at least $(1+\epsilon/3)d$. A simple union bound over the vertices in $N_{i,h}$ implies the following: for every $i = 0,\ldots,\ell$ we have that

$$\Pr\left[A_i \mid \mathcal{F}_S\right] \leq N_0 \exp\left(-\epsilon^2 d/27\right) \leq \exp\left(-\epsilon^2 d/30\right), \tag{24}$$

where N_0 is defined in (17). Also, it holds that

$$\Pr\left[B_i \mid \mathcal{F}_S\right] \leq \binom{N_0}{2}\frac{d}{n} \leq \frac{d^{5h}}{n}. \tag{25}$$

The above follows by noting B_i occurs, if there is an edge between the vertices $N_{i,h}$ which is not exposed during the revelation of the sets $\bigcup_{s=0}^{h} L_{i,s}$. The probability of having such an edge is upper bounded by the expected number of such edges.

Combining (24) and (25) with a simple union bound we get that

$$\Pr\left[A_i \cup B_i \mid \mathcal{F}_S\right] \leq \exp\left(-\epsilon^2 d/35\right). \tag{26}$$

Let R be the number of subgraphs $N_{i,h}$, for $i \in \{1,\ldots,\ell\}$, such that the event $A_i \cup B_i$ holds. Eq. (26) implies that for R we have the following: For any $x \in \{1,\ldots,\ell\}$ it holds that

$$\Pr[R = x] \leq \binom{\ell}{x} z_0^x (1-z_0)^{\ell-x}, \tag{27}$$

where $z_0 = \exp\left(-\epsilon^2 d/35\right)$. Also, we have that

$$\begin{aligned}
\Pr[F = s] &= \sum_{x=0}^{s} \Pr[R = x]\Pr[F = s \mid R = x] \\
&\leq \sum_{x=0}^{s} \binom{\ell}{x} z_0^x (1-z_0)^{\ell-x} \Pr[F = s \mid R = x] \\
&\leq \sum_{x=0}^{s} \binom{\ell}{x} z_0^x \cdot \Pr[F = s \mid R = x],
\end{aligned} \tag{28}$$

where the last inequality follows from the fact that $(1-z_0)^{\ell-x} \leq 1$.

We proceed by bounding appropriately the quantity $\Pr[F = s \mid R = x]$. For this, let Z be the number of pairs of subgraphs $N_{i,h}, N_{j,h}$ for which the event $E_{i,j}$ holds, for $i,j = 0,1,\ldots,\ell$. Given that $R = x$, so as to have $F = s$ there should be at least $\lceil (s-x)/2 \rceil$ pairs $N_{i,h}, N_{j,h}$ such that $E_{i,j}$ holds, i.e.

$$\Pr[F = s \mid R = x] \leq \Pr[Z \geq \lceil (s-x)/2 \rceil \mid R = x]. \tag{29}$$

Given some i and j, let J_1 be a subset of events $E_{i',j'}$ such that $E_{i,j} \notin J_1$. Also, let J_2 any subset of events $A_{i'}, B_{i'}$. Let $\mathcal{F}_J$ be the σ-algebra generated by the events in $J_1 \cup J_2$.

Noting that the expected number of edges between $N_{i,h}$ and $N_{j,h}$ is at most $N_0^2 d/n$, we have that

$$\Pr\left[E_{ij} \mid \mathcal{F}_J\right] \leq N_0^2 d/n \leq d^{5h}/n.$$

The above inequality implies that for any integer $x \geq 0$ and $z_1 = d^{5h}/n$, we have

$$\begin{aligned}
\Pr[Z \geq x] &\leq \sum_{r \geq x} \binom{\binom{\ell+1}{2}}{r} (z_1)^r (1-z_1)^{\binom{\ell+1}{2}-r} \\
&\leq \sum_{r \geq x} \binom{\binom{\ell+1}{2}}{r} (z_1)^r \leq \sum_{r \geq x} \left(\frac{(\ell+1)^2 e z_1}{2r} \right)^r \qquad \text{[since } \tbinom{n}{i} \leq (ne/i)^i\text{]} \\
&\leq 2 \left(\frac{(\ell+1)^2 e z_1}{2x} \right)^x \leq (4n^{-1} \log^4 n)^x, \qquad (30)
\end{aligned}$$

where the last inequality follows due to our assumption that $\ell \leq (\log n)^2$.

Plugging (30) , (29) into (28) we get that

$$\begin{aligned}
\Pr[F = s] &\leq \sum_{x=0}^{s} \binom{\ell}{x} z_0^x (4n^{-1} \log^4 n)^{(s-x)/2} \\
&\leq \sum_{x=0}^{s} \binom{\ell}{s-x} z_0^{s-x} (2n^{-1/2} \log^2 n)^x \\
&\leq \binom{\ell}{s} z_0^s \sum_{x=0}^{s} \binom{\ell}{s-x} \binom{\ell}{s}^{-1} [(2/z_0) n^{-1/2} \log^2 n]^x \\
&\leq \binom{\ell}{s} z_0^s \sum_{x=0}^{s} \frac{s!}{(s-x)!} \frac{(\ell-s)!}{(\ell-s+x)!} [(2/z_0) n^{-1/2} \log^2 n]^x \\
&\leq \binom{\ell}{s} z_0^s \sum_{x=0}^{s} \left(\frac{s}{\ell-s+1} \right)^x [(2/z_0) n^{-1/2} \log^2 n]^x \\
&\leq \binom{\ell}{s} z_0^s \frac{1}{1-n^{-2/5}},
\end{aligned}$$

where in the last inequality we use the fact that $s \leq \ell \leq (\log n)^2$ and $z_0 = \Theta(1)$. The lemma follows.

10 Proof of Proposition 2

For some vertex $w_j \in P$ such that $w_j \notin \boldsymbol{B}$ we have that $N_{j,h}$ is not `Fail`. That is, $N_{j,h}$ is a tree of maximum degree less than $(1+\epsilon/3)d$. For such $N_{j,h}$ we assume w_j to be the root.

If the height of $N_{j,h}$ is less than h, then no vertex in $N_{j,h}$ belongs to $\boldsymbol{B}$. For such tree, the proposition is trivially true. For the rest of the proof we assume that the height of $N_{j,h}$ is h.

From [12] we have the following theorem.

Theorem 9 (Jonasson 2001) *Let Δ, h be sufficiently large integers and let $k \geq \Delta + 2$. Let T be a complete Δ-ary tree of height h. Let r be the root and let L be the leaves of T. Also, let X be a random k-colouring of the tree. For any $c \in [k]$ it holds that*

$$\max_{\sigma \in \Omega_{T,k}} \left| \Pr[X(r) = c \mid X(L) = \sigma_L] - k^{-1} \right| \leq k^{-1} \phi_k(h),$$

where the quantity $\phi_k(h) \geq 0$ which tends to zero as $h \to \infty$.

Theorem 9 establishes the *Gibbs uniqueness* condition for the random colourings of a Δ-ary tree. In Proposition 3 we extend the previous result to trees of *maximum degree* Δ.

Proposition 3 *Let Δ, h be sufficiently large integers and $k \geq \Delta + 2$. Let T be a tree of height h and* maximum *degree at most Δ. Let r, L_0 denote the root and the vertices at level h, respectively. For X a random k-colouring of T, the following is true:*

For $\phi_k(h)$ as in Theorem 9 and for any $c \in [k]$ it holds that

$$\max_{\sigma \in \Omega_{T,k}} \left| \Pr[X(r) = c \mid X(L_0) = \sigma_{L_0}] - k^{-1} \right| \leq k^{-1}\phi_k(h).$$

The proof of Proposition 3 appears in Section 11.

Proof of Proposition 2: We let μ_N be the Gibbs distribution over the k-colourings of $\boldsymbol{N}$, while we let μ_{w_j} be the marginal of μ_N on $w_j \in P$. For $\sigma \in \Omega_{\boldsymbol{N},k}$ we let $t_\sigma \subseteq [k]$ contain all the colours that are used from σ to colour the vertices in Γ. It is elementary that $|t_\sigma| \leq 2$. Also, it holds that

$$\Pr[X(v_j) = c \mid N_{j,h}, X_\Gamma = \sigma_\Gamma] = (k - |t_\sigma|)^{-1}, \tag{31}$$

since we have assumed that $N_{j,h}$ is not `Fail`, the structure of $N_{j,h}$ is treelike. The above holds for any $N_{j,h} \in \mathcal{T}(d, \epsilon, h)$.

Let $\boldsymbol{N}'$ be the graph derived from $\boldsymbol{N}$ be deleting the edges of P which are incident to w_j. Let ν be the Gibbs distribution over the k-colourings of $\boldsymbol{N}'$, while let ν_{w_j} be the marginal of ν on w_j. For any $\sigma \in \Omega_{N,k}$ and any $c \in [k] \backslash t_\sigma$, let X be a random k-colouring of $\boldsymbol{N}$, then

$$\Pr[X(v_j) = c \mid N_{j,h}, X_{\boldsymbol{B}^+} = \sigma_{\boldsymbol{B}^+}] = \frac{\nu_{w_j}^{\sigma_{\boldsymbol{B}^+}}(c)}{1 - \nu_{w_j}^{\sigma_{\boldsymbol{B}^+}}(t_\sigma)}, \tag{32}$$

where $\nu_j^{\sigma_{\boldsymbol{B}^+}}(\cdot)$ denotes the distribution ν_j conditional that $\boldsymbol{B}^+$ is coloured $\sigma_{\boldsymbol{B}^+}$.

The proposition will follows by showing that the r.h.s. of (32) and (31) are sufficiently close. For this, we need to estimate $\nu_{w_j}^{\sigma_{\boldsymbol{B}^+}}(c)$. In particular, we show that for any $c \in [k]$ it holds that

$$\left| \nu_{w_j}^{\sigma_{\boldsymbol{B}^+}}(c) - k^{-1} \right| \leq k^{-1} \cdot \phi_k(h), \tag{33}$$

where $\phi_k(h) : \mathbb{N}^+ \to \mathbb{R}_{\geq 0}$ is the function defined in Theorem 9.

In the graph $\boldsymbol{N}'$, the component of w_j, i.e. $N_{j,h}$ is a tree and it is only the vertices at distance h from w_j that belong to $\boldsymbol{B}$. The colouring of the vertices in Γ does not affect the colour assignment of w_j, since we have deleted the edges of P which are incident to w_j. Since $N_{j,h} \in \mathcal{T}(d, \epsilon, h)$, Proposition 3 implies that (33) is indeed true for any $N_{j,h} \in \mathcal{T}(d, \epsilon, h)$.

Combining (33) and (32) we get that

$$\left| \Pr[X(v_j) = c \mid N_{j,h}, X_{\boldsymbol{B}^+} = \sigma_{\boldsymbol{B}^+}] - (k - |t_\sigma|)^{-1} \right| \leq 10k^{-1}\phi_k(h). \tag{34}$$

The proposition follows from (34) and (31) and setting $f_\epsilon(h) = 10\phi_k(h)$. ◊

11 Proof of Proposition 3

Let T' be a supertree of T such that T' is a complete Δ-ary tree of height h. That is, T and T' have the same height. Also, both trees have the same root r. We denote with L the set of vertices at level h in T'. $L_0 \subseteq L$ is the set of vertices which are at level h in both T and T'.

For T and T' we have the following result.

Lemma 6 *Assume that $k \geq \Delta + 2$. Let X, Y be random k-colourings of T, T', respectively. Also, let σ be any k-colouring of T. For any $c \in [k]$ it holds that*

$$\Pr\left[X(r) = c \mid X(L_0) = \sigma_{L_0}\right] = \Pr\left[Y(r) = c \mid Y(L_0) = \sigma_{L_0}\right].$$

The proof of Lemma 6 appears in Section 11.1.

Given Lemma 6, we show the proposition by working as follows: Let X, Y be a random k-colouring of T and T', respectively. Let $\tau \in \Omega_{T,k}$ be such that τ_{L_0} maximizes the following quantity,

$$|\Pr[X(r) = c \mid X(L_0) = \tau_{L_0}] - k^{-1}|.$$

By Lemma 6, we have that $\Pr[X(r) = c \mid X(L_0) = \tau_{L_0}] = \Pr[Y(r) = c \mid Y(L_0) = \tau_{L_0}]$. It holds that

$$|\Pr[X(r) = c \mid X(L_0) = \tau_{L_0}] - k^{-1}| \leq \max_{\sigma \in \Omega_{T',k}} |\Pr[Y(r) = c \mid Y(L_0) = \sigma_{L_0}] - k^{-1}|,$$

where σ varies over all the proper colourings of T'. The proposition follows by using Theorem 9 to bound the r.h.s. of the inequality above.

11.1 Proof of Lemma 6

For the tree T (resp. the tree T') and a vertex v, let T_v (resp. T'_v) denote the subtree that contains the vertex v once we delete the edge of T (resp. T') that connects v and its parent. For the tree T_v (resp. T'_v) the root is the vertex v.

Consider the random colourings X, Y of the trees T and T', respectively, with boundary condition σ_{L_0}. Also, consider the following random variables: For every vertex $v \in T$, (resp. T') we consider the subtree T_v (resp. T'_v) and the random colouring X^v (resp. Y^v) on this tree, with boundary conditions set as follows: Letting $L_v = L_0 \cap T_v$, then the boundary condition for both X^v and Y^v is σ_{L_v}.

We denote with C the set of the children of the root r which belong to both trees, T, T'. Also, we denote with S be the set of children of r which belong only to the tree T'.

The proof is by induction on the height of the tree h. We start with $h = 1$. Since the height of the tree is 1, it holds that $C = L_0$. Clearly for any color which appears in the boundary it holds that neither X nor Y is going to use it for colouring the root. Let $U \subset [k]$ contain all the colours that are not used by the boundary condition σ_{L_0}. For any $c \in U$ it holds that

$$\begin{aligned}\Pr[Y(r) = c \mid Y(L_0) = \sigma_{L_0}] &= \frac{\prod_{v\in S}(1 - \Pr[Y^v(v) = c]) \times \prod_{v\in C}(1 - \Pr[Y^v(v) = c])}{\sum_{q\in[k]} \left(\prod_{v\in S}(1 - \Pr[Y^v(v) = q]) \times \prod_{v\in C}(1 - \Pr[Y^v(v) = q]\right)} \\ &= \frac{\prod_{v\in S}(1 - \Pr[Y^v(v) = c])}{\sum_{q\in U} \prod_{v\in S}(1 - \Pr[Y^v(v) = q])}.\end{aligned}$$

To see why the second inequality holds consider the following: If $q \notin U$, then we have that $\prod_{v\in C}(1 - \Pr[Y^v(v) = q]) = 0$, since, we have assumed that there is $v \in C$ such that $Pr[Y^v(v) = q] = 1$. On the other hand, if $q \in U$, then $\prod_{v\in C}(1 - \Pr[Y^v(v) = q]) = 1$ since, by definition, for every $v \in C$ it holds that $\Pr[Y^v(v) = q] = 0$. Furthermore, it is direct that

$$\Pr[Y(r) = c \mid Y(L_0) = \sigma_{L_0}] = \frac{(1 - 1/k)^{|S|}}{|U|(1 - 1/k)^{|S|}} = \frac{1}{|U|} = \Pr[X(r) = c \mid X(L_0) = \sigma_{L_0}].$$

Assume now that our hypothesis is true for trees of height $h - 1$, for some $h \geq 2$. We are going to show that the hypothesis is true for trees of height h, too. It holds that

$$\begin{aligned}\Pr[X(r) = c \mid X(L_0) = \sigma_{L_0}] &= \frac{\prod_{v\in C}(1 - \Pr[X^v(v) = c])}{\sum_{q\in[k]} \prod_{v\in C}(1 - \Pr[X^v(v) = q])} \\ &= \frac{\prod_{v\in C}(1 - \Pr[Y^v(v) = c])}{\sum_{q\in[k]} \prod_{v\in C}(1 - \Pr[Y^v(v) = q])},\end{aligned} \tag{35}$$

where the second equality follows from the induction hypothesis. Also, it holds that

$$\begin{aligned}\Pr[Y(r)=c \mid Y(L_0)=\sigma_{L_0}] &= \frac{\prod_{v\in S}(1-\Pr[Y^v(v)=c])\times\prod_{v\in C}(1-\Pr[Y^v(v)=c])}{\sum_{q\in[k]}\left(\prod_{v\in S}(1-\Pr[Y^v(v)=q])\times\prod_{v\in C}(1-\Pr[Y^v(v)=q]\right)}\\ &= \frac{(1-1/k)^{|S|}\prod_{v\in C}(1-\Pr[Y^v(v)=c])}{\sum_{q\in[k]}\left((1-1/k)^{|S|}\prod_{v\in C}(1-\Pr[Y^v(v)=q]\right)}\\ &= \frac{\prod_{v\in C}(1-\Pr[Y^v(v)=c])}{\sum_{q\in[k]}\prod_{v\in C}(1-\Pr[Y^v(v)=q])}, \end{aligned} \tag{36}$$

where the second equality holds because for every $v\in S$ it holds $\Pr[Y^v(v)=c]=k^{-1}$. Observe that if $v\in S$, then the subtree T'_v contains no vertex u which also belongs to T, thus Y^v has no boundary conditions at all. The lemma follows from (35) and (36).

12 Proof of Theorem 5

For proving Theorem 5 we need the following result.

Lemma 7 *For any $c,q\in[k]$ such that $c\neq q$, it holds that* $\texttt{Switching}(G,v,\cdot,q):S_q(c,c)\to S_c(q,c)$ *is a* bijection.

Proof: For any $\sigma\in S_q(c,c)$, it holds that $\texttt{Switching}(G,v,\sigma,q)\in S_c(q,c)$. This follows from Lemma 1 and the definition of the sets $S_q(c,c)$ and $S_c(q,c)$.

It suffices to show that the mapping $\texttt{Switching}(G,v,\cdot,q):S_q(c,c)\to S_c(q,c)$ it is *one-to-one* and it is *surjective*, i.e. it has range $S_c(q,c)$. For showing both properties we use the following observation: If for some $\tau\in S_c(q,c)$ and $\xi\in S_q(c,c)$ it holds that $\tau=\texttt{Switching}(G,v,\xi,q)$, then it also holds that $\xi=\texttt{Switching}(G,v,\tau,c)$.

As far as surjectiveness is regarded, it suffices to have that for every $\tau\in S_c(q,c)$ there exists $\xi\in S_q(c,c)$ such that $\texttt{Switching}(G,v,\xi,q)=\tau$. From the above observation we get that each $\tau\in S_c(q,c)$ is the image of $\xi\in S_q(c,c)$ for which it holds that $\xi=\texttt{Switching}(G,v,\tau,c)$. Furthermore, we observe that this ξ is unique. This implies that $\texttt{Switching}(G,v,\cdot,q)$ is *one-to-one*, too.

The lemma follows. ◊

Proof of Theorem 5: Let X,Y be the input and the output of `Update`, respectively. X is distributed uniformly at random among the k-colourings of G. Also, let Z be a random variable distributed as in ν, the uniform distribution over the good k-colourings of G.

The theorem will follow by providing a coupling of Z and Y such that

$$\Pr[Z\neq Y]\leq\alpha.$$

First, we need the following observations: For any $q,c\in[k]$ such that $c\neq q$, it holds that

$$\Pr[Z(v)=q\mid Z(u)=c]=\Pr[X(v)=q\mid X(u)=c,X(v)\neq c]=(k-1)^{-1} \tag{37}$$

and

$$\Pr[X(v)=X(u)=c\mid X\text{ is bad }]=k^{-1}. \tag{38}$$

All the above equalities follow due to symmetry between the colours. Also, it is direct to show that

$$\Pr[Y(v)=q\mid X(u)=c]=(k-1)^{-1}. \tag{39}$$

In particular, (39) holds because $Y(v)$ is set according to the following rules: if X is good, then we have that $X = Y$ and (37) holds. On the other hand, if X is bad and $X(u) = c$, then $Y(v)$ is chosen uniformly at random from $[k]\backslash\{c\}$.

Now we are going to describe the coupling. We need to involve the variable X in the coupling, since Y depends on it. At the beginning, we set $Z(u) = X(u)$, also we set $Z(v) = Y(v)$. From (37), (38) and (39), it is direct that $Z(u)$ and $Z(v)$ are set according to the appropriate distribution.

We need to consider two cases, depending on whether X is a good or a bad colouring. For each case we have different couplings. Then it holds that

$$\Pr[Y \neq Z] \leq \Pr[Y \neq Z \mid X \text{ is good}] + \Pr[Y \neq Z \mid X \text{ is bad}]. \tag{40}$$

If X is good, then it is distributed uniformly at random among the good colourings of G. That is, X and Z are identically distributed. That is, if X is good, then there is a coupling such that $X = Z$ with probability 1. Also, from `Update` we have that $X = Y$. It is direct that if X is good, then there is a coupling such that

$$\Pr[Y \neq Z \mid X \text{ is good}] = 0. \tag{41}$$

On the other hand, if X is a bad colouring, the situation is as follows: If $X(u) = X(v) = c$, for some $c \in [k]$, then $Z(u) = c$ and $Z(v) = q$ for some $q \in [k]\backslash\{c\}$ and $Y(v) = q$. We let the event $E_{c,q} =$ "$X(u) = X(v) = Z(u) = c$ and $Y(v) = Z(v) = q$ while $X \in S_q(c, c)$ and $Z \in S_c(q, c)$". Also, let the event $E = \bigcup_{c,q\in[k]:c\neq q} E_{c,q}$.

In the coupling we are distinguishing the cases where the event E occurs from those that is does not. For each case we have different couplings. It holds that

$$\Pr[Y \neq Z \mid X \text{ is bad}] \leq \Pr[Y \neq Z \mid E, X \text{ is bad}] + \Pr[\bar{E} \mid X \text{ is bad}], \tag{42}$$

where $\bar{E}$ is the complement of E. The theorem follows by showing that the r.h.s. of (42) is at most α. From the definition of the quantity α (Definition 5), it holds that

$$\Pr[X \in S_q(c, c) \mid X(u) = X(v) = c] \geq 1 - \alpha,$$

also, it holds that

$$\Pr[Z \in S_c(q, c) \mid Z(u) = c, Z(v) = q] \geq 1 - \alpha,$$

for any $c, q \in [k]$ and $q \neq c$. The above implies that, when X is bad, there is a coupling such that

$$\Pr[E \mid X \text{ is bad}] \geq 1 - \alpha. \tag{43}$$

It remains to describe a coupling of Z, Y, when X is bad and E occurs (i.e. bound $\Pr[Y \neq Z \mid E, X \text{ is bad}]$). For this, we need the following claim.

Claim 3 *Conditional on the event $E_{c,q}$, Y is distributed uniformly over $S_c(q, c)$.*

Proof: From Lemma 7 we have that $\texttt{Switching}(G, v, \cdot, q) : S_q(c, c) \to S_c(q, c)$ is a bijection. The existence of this bijection implies that $|S_q(c, c)| = |S_c(q, c)|$. Also, for each $\tau \in S_c(q, c)$ there is a unique $\xi \in S_q(c, c)$ such that $\texttt{Switching}(G, v, \xi, q) = \tau$. Clearly $\Pr[Y = \tau \mid E_{c,q}] = \Pr[X = \xi \mid E_{c,q}]$.

Conditional on the event $E_{c,q}$, the random variable X is distributed uniformly over $S_q(c, c)$. Thus, $\Pr[Y = \tau | E_{c,q}] = |S_q(c, c)|^{-1} = |S_c(q, c)|^{-1}$, for any $\tau \in S_c(q, c)$. The claim follows. $\Diamond$

It is direct that conditional on $E_{c,q}$ the random variable Z is distributed uniformly at random in $S_c(q, c)$. Also, observe that conditional on that X is bad and E occurring, we are going to have $Z(v) = Y(v)$ and $Z(u) = Y(u)$. All these imply that there is a coupling of Z, Y such that

$$\Pr[Y \neq Z \mid X \text{ is bad}, E] = 0. \tag{44}$$

Plugging (43) and (44) into (42), we get that

$$\Pr[Y \neq Z \mid X \text{ is bad}] \leq \alpha.$$

The theorem follows by plugging the above bound and (41) into (40). $\Diamond$

13 The rest of the proofs

13.1 Lemma 1

We show that for any $\sigma \in \Omega_{G,k}$, it holds that $\texttt{Switching}(G, v, \sigma, q)$ returns a proper colouring of G. Assume the contrary, i.e. there is $\sigma \in \Omega_{G,k}$ such that for $\tau = \texttt{Switching}(G, v, \sigma, q)$ it holds that $\tau \notin \Omega_{G,k}$.

Let the disagreement graph $\boldsymbol{Q} = \boldsymbol{Q}(G, v, \sigma, q)$. Since τ is non-proper is has at least one monochromatic edge. The monochromatic edge can be incident either to two vertices in $\boldsymbol{Q}$ or to some vertex in $\boldsymbol{Q}$ and some vertex outside $\boldsymbol{Q}$. We are going to show that neither of the two cases can happen.

$\texttt{Switching}(G, v, \sigma, q)$ cannot create any monochromatic edge between two vertices in $\boldsymbol{Q}$. To see this, note that the disagreement graph $\boldsymbol{Q}$ is bipartite and σ specifies exactly one colour for each part of the graph. $\texttt{Switching}(G, v, \sigma, q)$ just exchanges the colours of the two parts in the graph. Clearly this operation cannot generate a monochromatic of the first kind.

$\texttt{Switching}(G, v, \sigma, q)$ cannot cause any monochromatic edge between a vertex in $\boldsymbol{Q}$ and some vertex outside $\boldsymbol{Q}$, either. This follows by the fact that the disagreement graph is maximal. That is, there is no vertex w outside $\boldsymbol{Q}$ such that $\sigma_w \in \{q, c\}$ while at the same time w is adjacent to some vertex in $\boldsymbol{Q}$. Since the recolouring that $\texttt{Switching}(G, v, \sigma, q)$ does, involves only vertices coloured c, q, no monochromatic edge of the second kind can be generated, too.

The lemma follows.

13.2 Lemma 2

The time complexity of computing $\texttt{Switching}(G, v, \sigma, q)$ is dominated by the time we need to reveal the disagreement graph $\boldsymbol{Q} = \boldsymbol{Q}(G, v, \sigma, q)$. We will show that we need $O(|E(G)|)$ steps to get $\boldsymbol{Q}$.

We reveal the graph $\boldsymbol{Q}$ in steps $j = 0, \ldots, |E(G)|$. At step 0, we have $\boldsymbol{Q}(0)$ which contains only the vertex v. Given $\boldsymbol{Q}(j)$ we construct $\boldsymbol{Q}(j+1)$ as follows: Pick some edge which is incident to a vertex in $\boldsymbol{Q}(j)$. If the other end of this edge is incident to a vertex outside $\boldsymbol{Q}(j)$ that is coloured either σ_v or q, then we get $\boldsymbol{Q}(j+1)$ by inserting this edge and the vertex into $\boldsymbol{Q}(j)$. Otherwise $\boldsymbol{Q}(j+1)$ is the same as $\boldsymbol{Q}(j)$. We never pick the same edge twice in the process above.

The lemma follows by noting that the process has at most $|E|$ steps, while at the end we get $\boldsymbol{Q}$.

13.3 Theorem 7

For $i = 0, \ldots, r$ consider the following: Let μ_i denote the uniform distribution over the k-colourings of G_i. Also let $\hat{\mu}_i$ denote the distribution of Y_i, where Y_i is the colouring that the algorithm assigns to the graph G_i. Finally, let ν_i denote the distribution of the output colouring of $\texttt{Update}(G_i, v_i, u_i, X_i, k)$ where X_i is distributed as in μ_i.

The theorem follows by showing that that

$$||\mu_r - \hat{\mu}_r|| \leq \sum_{i=0}^{r-1} \alpha_i. \tag{45}$$

Theorem 5 implies the following: For every $i = 1, \ldots, r$ it holds that

$$||\mu_i - \nu_{i-1}|| \leq \alpha_{i-1}, \tag{46}$$

It suffices to show that

$$||\mu_r - \hat{\mu}_r|| \leq \sum_{i=1}^{r} ||\mu_i - \nu_{i-1}||, \tag{47}$$

since it is direct that (45) follows from (46) and (47).

For getting (47), we are going to show for any $i = 1, \ldots, r$ the following is true:

$$||\nu_{i-1} - \hat{\mu}_i|| \leq ||\mu_{i-1} - \hat{\mu}_{i-1}||. \tag{48}$$

From (48) we get to (47) by working as follows: Using the triangle inequality, we have that

$$\begin{aligned} ||\mu_r - \hat{\mu}_r|| &\leq ||\mu_r - \nu_{r-1}|| + ||\nu_{r-1} - \hat{\mu}_r|| \\ &\leq ||\mu_r - \nu_{r-1}|| + ||\mu_{r-1} - \hat{\mu}_{r-1}||. \qquad \text{[from (48)]} \end{aligned}$$

We work with the term $||\mu_{r-1} - \hat{\mu}_{r-1}||$, above, in the same way as we did with $||\mu_r - \hat{\mu}_r||$ and so on. This sequence of substitutions and the fact that $||\mu_0 - \hat{\mu}_0|| = 0$, yield (47).

It remains to show (48). For this, let X_{i-1} be a random k-colouring of the graph G_{i-1} and let $Z_i = \texttt{Update}(G_{i-1}, v_{i-1}, u_{i-1}, X_{i-1}, k)$. It is direct that Z_i is distributed as in ν_{i-1}. Let Y_{i-1}, Y_i be the colouring that the algorithm assigns to the graphs G_{i-1}, G_i, respectively. Clearly it holds that $Y_i = \texttt{Update}(G_{i-1}, v_{i-1}, u_{i-1}, Y_{i-1}, k)$

So as to bound $||\nu_{i-1} - \hat{\mu}_i||$ we consider the following coupling of Z_i and Y_i: We couple X_{i-1} and Y_{i-1} *optimally*. Then from X_{i-1} and Y_{i-1}, we get Z_i and Y_i, respectively, as described above. By the coupling lemma we have the following

$$||\nu_{i-1} - \hat{\mu}_i|| \leq \Pr[Z_i \neq Y_i] \leq \Pr[Z_i \neq Y_i \mid X_{i-1} = Y_{i-1}] + \Pr[X_{i-1} \neq Y_{i-1}]. \tag{49}$$

It is direct that if $X_{i-1} = Y_{i-1}$, then there is a coupling which yield $Z_i = Y_i$ with probability 1. That is, $\Pr[Z_i \neq Y_i \mid X_{i-1} = Y_{i-1}] = 0$. Also, since we have coupled X_{i-1} and Y_{i-1} optimally, it holds that

$$\Pr[X_{i-1} \neq Y_{i-1}] = ||\mu_{i-1} - \hat{\mu}_{i-1}||. \tag{50}$$

Plugging (50) into (49) and using the fact that $\Pr[Z_i \neq Y_i \mid X_{i-1} = Y_{i-1}] = 0$, we get (48). The theorem follows.

13.4 Lemma 3

It suffices to show that with probability at least $1 - n^{-2/3}$ for any two cycles in $\boldsymbol{G}$, of maximum length $(\log_d n)/9$ do not share edges and vertices with each other. Assume the opposite, i.e. that there are at least two such cycles that intersect with each other. Then, there must exist a subgraph of $\boldsymbol{G}$ that contains at most $(2/9)\log_d n$ vertices while the number of edges exceeds by 1, or more, the number of vertices.

Let D be the event that in $\boldsymbol{G}$ there exists a set of r vertices which have $r+1$ edges between them, for $r \leq (2\log_d n)/9$. The lemma follows by showing that $\Pr[D] \leq n^{-2/3}$.

We have the following:

$$\begin{aligned} \Pr[D] &\leq \sum_{r=1}^{(2/9)\log_d n} \binom{n}{r} \binom{\binom{r}{2}}{r+1} (d/n)^{r+1} (1-d/n)^{\binom{r}{2}-(r+1)} \\ &\leq \sum_{r=1}^{(2/9)\log_d n} \left(\frac{ne}{r}\right)^r \left(\frac{r^2 e}{2(r+1)}\right)^{r+1} (d/n)^{r+1} \leq \frac{e \cdot d}{2n} \sum_{r=1}^{(2/9)\log_d n} r \left(\frac{e^2 d}{2}\right)^r \\ &\leq \frac{C \log n}{n} \left(\frac{e^2 d}{2}\right)^{(2/9)\log_d n}. \end{aligned}$$

Let $\gamma = \frac{2\log(e^2 d/2)}{9\log d}$. The quantity in the r.h.s. of the last inequality, above, is of order $\Theta(n^{\gamma-1}\log n)$. Taking large d it holds that $\gamma < 0.25$. Consequently, we get that $\Pr[D] \leq n^{-2/3}$. The lemma follows.

Acknowledgement. The author of this work would like to thank Amin Coja-Oghlan and Elchanan Mossel for the fruitful discussions we had. Also, I would like to thank the anonymous reviewers for helping me improved the content of the paper.

References

[1] D. Achlioptas and A. Coja-Oghlan, *Algorithmic Barriers from Phase Transitions.* In proc. of FOCS 2008: 793-802.

[2] D. Aldous. *Random walks of finite groups and rapidly mixing Markov chains.* In: Séminaire de Probabilités XVII 1981/82, Springer-Verlag, Berlin. pp. 243-297.

[3] J. van den Berg and C. Maes. *Disagreement percolation in the study of Markov fields.* Annals of Probability **22**, (1994) pp. 794-763.

[4] A. Braunstein, M. Mézard, R. Zecchina: *Survey propagation: an algorithm for satisfiability.* Random Structures and Algorithms **27** (2005) 201-226

[5] M. Dyer, A. Flaxman, A. M. Frieze and E. Vigoda. *Random colouring sparse random graphs with fewer colours than the maximum degree*. Random Struct. and Algorithms **29**, pp. 450-465, 2006.

[6] M. Dyer, A. M. Frieze, A. Hayes, E. Vigoda. *Randomly colouring constant degree graphs.* In proc. of 45th FOCS, pp 582-589, 2004.

[7] C. Efthymiou. *MCMC sampling colourings and independent sets of G(n, d/n) near uniqueness threshold.* In Proc. of 25th ACM-SIAM, Symposium on Discrete Algorithms (SODA'14), pp 305-316, 2014.

[8] C. Efthymiou, *A simple algorithm for random colouring $G(n, d/n)$ using $(2+\epsilon)d$ colours.* In Proc. of 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA'12) pp 272-280, 2012.

[9] C. Efthymiou, *Switching colouring of $G(n, d/n)$ for sampling up to Gibbs Uniqueness Threshold.* In proc of 22nd European Symposium on Algorithms (ESA), pp 371-381, 2014.

[10] A. M. Frieze, J. Vera. *On randomly colouring locally sparse graphs.* Discrete Mathematics & Theoretical Computer Science **8** (1), pp 121-128 2006

[11] L. A. Goldberg, R. A. Martin, M. Paterson. *Strong Spatial Mixing with Fewer Colors for Lattice Graphs.* SIAM J. Comput. **35** (2), pp 486-517 2005

[12] J. Jonasson. *Uniqueness of Uniform Random Colorings of Regular Trees* In Statistics & Probability Letters **57**, pp 243–248, 2001.

[13] S. Janson, T. Luczak and A. Ruciński. *Random graphs.* Wiley and Sons, Inc. 2000.

[14] M. Jerrum and A. Sinclair. *The Markov chain Monte Carlo method:an approach to approximate counting and integration.* In Approximation Algorithms for NP-hard problems (Dorit, Hochbaum ed.) PWS 1996.

[15] F. Martinelli, A. Sinclair and D. Weitz. Fast mixing for independent sets, colorings and other models on trees. In proc. of 15th SODA, pp 456-465, 2004.

[16] M. Molloy. *The freezing threshold for k-colourings of a random graph.* In proc. of the 44th ACM Symposium on Theory of Computing (STOC'12), pp 921-930, 2012.

[17] E. Mossel and A. Sly. *Gibbs Rapidly Samples Colorings of* $G_{n,d/n}$. In Probability Theory and Related Fields, **148**, No 1-2, 2010.

[18] E. Vigoda. *Improved bounds for sampling colorings.* Journal of Mathematical Physics, **41** (3), pp 1555-1569, 2000. A preliminary version appears in FOCS 1999.

[19] M. J. Wainright and M. Jordan. *Graphical Models, Exponential Families and Variational Inference.* In Foundations and Trends in Machine Learning **1** (1-2), 2008.

[20] D. Weitz. *Counting independent sets up to the tree threshold* In the proc.of 38th Annual ACM Symposium on Theory of Computing (STOC) 2006, pp 140–149.

[21] Y. Yin, C. Zhang, *Sampling colorings almost uniformly in sparse random graphs*. In http://arxiv.org/abs/1503.03351

[22] L. Zdeborová, F. Krzakala. *Phase transition in the Coloring of Random Graphs*. Phys. Rev. E **76**, 031131(2007). preprint: arxiv:0704.1269 .